NATIONAL COMPUTATIONAL INFRASTRUCTURE

Providing Australian researchers with world-class high-end computing services

PROGRESS REPORT

July 2008 – June 2009

September 2009

Table of Contents

1. INTRODUCTION

This report outlines the progress and achievements of the National Computational Infrastructure (NCI) program during the twelve month period July 2008 – June 2009, and also alludes to ongoing progress during the period July – September 2009 that will be reported in detail in next year's Progress Report.

The NCI Project is funded as part of the Platforms for Collaboration (PfC) Investment Plan. Its purpose is to provide state-of-the-art national computational infrastructure facilities and services to the Australian research community and builds on the successful National Facility (NF) program of the Australian Partnership for Advanced Computing (APAC).

The NCI project was established and funded through a contract between The Australian National University (ANU) and the Commonwealth of Australia represented, at the time of signing (22 June 2007), by the Department of Education, Science and Training, and subsequently by the Department of Innovation, Industry, Science and Research (Department of Innovation) following changes to the structure of the Commonwealth Government following the November 2007 election.

The program has three main activities — Operational, Outreach and Policy/Planning — the status of which are summarised in Section 2 and are expanded upon in Sections 3 and 4 of this report.

More information concerning NCI can be found at its websites (www.nci.org.au and nf.nci.org.au).

2. PROJECT STATUS

The achievements of the NCI programs and related activities over the period July 2008-June 2009 are summarised below and detailed in the report on project activities (Sec. 3). This and the following section review progress against the Business Plan and comment on any variations from this.

Governance and Management

Steering Committee

The NCI Steering Committee was established in late 2007 during the implementation phase of NCI, with Professor Mark Wainwright appointed as the independent Chair of the NCI Steering Committee. Professor Jim Williams of ANU served as Acting Director from July 2007 – April 2008 prior to the substantive Director taking up his appointment in May 2008.

The initial governance structure was based on a partnership (co-investment) model in which the large investors (termed partners) would be members of the Steering Committee in their own right, while the smaller investors (termed affiliates) would select one of their number to represent them on the Steering Committee. Since its establishment in late 2007, the Steering Committee has had representation from ANU as the contracting organisation, together with CSIRO, the Bureau of Meteorology, Geoscience Australia, and the research intensive universities, all of which are leading research organisations having a commitment to, and a direct interest in, the advancement of Australia's high-end computing infrastructure, and which were deemed to be potential partners (co-investors) of NCI. This initial composition was agreed to by the Department of Innovation until either 30 September 2008, or until a potential partner became a actual partner through the appropriate level of co-investment.

During the course of the current reporting period, the Steering Committee evolved its thinking about the governance of NCI, particularly with respect to the framework that was based exclusively on co-investment / partnership, and also about ways in which to most effectively advance the national planning and policy agenda. In the previous annual report (September 2008), NCI foreshadowed the bifurcation of the strategic planning and operational roles of the NCI Steering Committee, proposing the establishment of a national council as an outgrowth of the NCI Steering Committee to promote, plan and coordinate the advancement of high-end computing in Australia, together with an operational committee, drawn from the partner investors, to oversee the NCI program.

The Department of Innovation on advice from AeRIC chose to defer further consideration of the concept of a national council, noting in its November 2008 reply "that the initiative may possibly be overtaken by other events such as the consideration of HPC needs for climate change and the Government's response to the Review of the National Innovation System". In order to provide continuity in the governance of NCI, the terms of membership for the Steering Committee were extended by the Department until 30 June 2009.

The subsequent signing of the NCI Partner Service Agreement by ANU and CSIRO on 24 December 2008 (each committing to annual co-investments exceeding $3M) confirmed CSIRO as a full partner under the existing framework.

Subsequently in the 2009-10 Business Plan submitted to the Department of Innovation on 31 March 2009, NCI proposed the formalisation and confirmation of the composition of the Steering Committee to provide a framework, both for overseeing the operations of NCI, as well as establishing a platform for the strategic planning of national high-end computing services. The proposed composition of the revised Steering Committee comprised the existing membership augmented by the University of Melbourne representing the interests of the Victorian Life Sciences Computation Initiative, established during the course of the reporting period. Prior to this, the University of Melbourne had been invited to provide an observer to Steering Committee meetings throughout the reporting year 2008-09.

This revised composition, which is listed in Appendix 2, was subsequently agreed to by the Department and was formalised through a Deed of Variation to the Funding Agreement signed in late May 2009.

The role of Steering Committee in national planning has been since clarified during July 2009, with the Department indicating its intention to refer such matters to the recently constituted National Research Infrastructure Council (NRIC).

Director, NCI

NCI commenced operations with Professor Jim Williams serving as Acting Director. The decision to advertise for a substantive Director was taken in late 2007, with the position advertised in January 2008. Following interviews in March 2008, Professor Lindsay Botten took up the role in late May 2008. While the late appointment of a substantive Director has caused some delays, these have been largely recovered and the program is well on track.

Planning, Policy and Access

Refinement of the Partnership Model

Substantial work on the development of a Resource Allocation Model (i.e., a shares model) continued during this reporting period and culminated in the development of the NCI Partner Service Agreement (PSA) which provides the framework by which co-investing partners engage with, and are provided with services by, NCI. The present partner model is framed in terms of long-term commitments to NCI and assists greatly in managing the cash flow. Organisations or consortia which invest in excess of $50,000 p.a. for the lifetime of the PSA (which extends until the end of 2011) are entitled to partner membership. Partner membership weights the value of partner contributions relative to NCRIS contributions (which provide for the Merit Allocation Scheme share) by factor of 1.1. In this way, Commonwealth funds leverage new partner investment, leading to a "win-win" situation in which both MAS and partners have shares on a significantly larger facility that could have been acquired solely through Commonwealth resources.

The signing of the PSA by ANU and CSIRO confirmed $6.7M p.a. of co-investment in NCI, relative to the annualised contribution from the Commonwealth of $6.5M. The subsequent signing of a number of smaller partners has increased this to almost $6.9M annually. From the beginning of 2010, this will increase by a further $560K p.a. (annualised over a four year period) through the Intersect Consortium of NSW joining as a partner.

Merit Allocation

The Merit Allocation Scheme, the mechanism by which that share of the facilities provided by NCRIS funds, operated under new processes that were foreshadowed in the last year's report. In summary, these revised processes increased the rigour of the assessment process without increasing the work required of researchers in preparing the application—by incorporating relevant information that researchers would have already prepared for applications to major funding bodies (e.g., ARC, NHMRC). In addition to this, the online submission of the applications was improved and overall the new processes were well received by researchers. The additional information allowed closer scrutiny of the research impact and, to that end, the Committee considered publication citation data more closely than in previous years.

The Merit Allocation Committee continues as a matrix which balances organisational/geographical considerations with the necessary span of discipline expertise. It continues its strong expert tradition, with all members having substantial personal records, and additionally, in the case of university researchers, strong portfolios of current research funding. One member of the committee is a Federation Fellow, while two others are ARC Professorial Fellows. There was some turnover in committee membership, with two long serving university members stepping down and being replaced, and, for the first time, CSIRO providing two representatives. The presence of two additional members assisted in balancing the ever increasing workloads and allowed for the inclusion of dedicated expertise in climate science for the first time.

The share allocated to the MAS for 2009 was 45.99%, down from that allocated for the previous year (73%) and returning to a level comparable to that available under APAC during its second phase of operations (2004-06). In this regard, the previous year (2008) was an anomaly in which ANU was the only substantial partner during the transition from APAC to NCI.

For the 2009 round of merit allocation, the total number of hours available to the MAS was 12.8M, the same as in 2008. With an increase in requests rising from 19.5M hours for 2008 through to 28.2M for 2009, the oversubscription ratio rose from 1.5 to 2.2. Fortuitously, however, the introduction of the new SGI Altix XE early in 2009 alleviated this problem, at least to some extent, since the new system exhibited a superior performance to that available on the existing peak system (SGI Altix 3700) on a range of codes (i.e., providing more useful computation in a unit of time than could be achieved on the older system). A comprehensive report on outcomes is provided in Sec. 3.1.1.1 and Appendices 3 and 4.

Discussions at the 2009 meeting of the Merit Allocation Committee have led to the consideration of a Flagship MAS intended particularly to provide the needs of the high impact research projects and the priorities of research communities, research organisations and also government funding agencies. The policy development work that is underway within NCI is outlined in Sec 3.2.2 — work which will be of real significance if Australia is to make the most effective use of its rapidly evolving high-end computing services. Also discussed, in the context of the evolving national environment, is the need for a National Merit Allocation scheme—a mechanism which NCI is uniquely well placed to host.

Outreach

Good progress has been made over the past year in the Outreach program, the major focus of which is in raising the profile and uptake of high-end computing in Australia. Over the past year, NCI has increased its engagement with research communities and leading research institutions to assist them in enhancing the impact and scale of Australian research achievements through high-end computing services.

Of particular relevance has been the work with Australia's climate community leading to the joint procurement (by ANU and the Bureau of Meteorology) of new interoperational HPC facilities. NCI has also contributed to the development of the National Framework for Australian Climate Change Science which ultimately informed allocations made under the new Super Science Initiative. On other fronts, the development of support through the Computational Tools and Techniques program for the climate science community is

being well received, as is the initiation of planning for the next generation system for climate change modelling through the involvement of a broad cross-section of the leadership of the climate science, impact and adaptation community.

Also of significance is the growing level of partner investment in NCI, with the CSIRO's confirmation of partnership and Intersect's intention to join the NCI partnership from early 2010 being important milestones. A fuller treatment of these and similar matters is provided in Sec. 3.2.2.

The current reporting year has also seen considerable activity in the development of new promotional materials for NCI. At the time of writing, NCI is planning the redevelopment of its website to combine the existing site and the National Facility website into a single coherent entity which incorporates a significant focus on the research outcomes and their impact. In preparation for this, and also the launch of the new National Facility, considerable work has been devoted to the development of a gallery of research highlights that feature the work of leading researchers. In the first phase of this, a number of projects which are not only of high impact but which are also "media friendly" have been identified and developed with the assistance of a leading science journalist. The materials that have been prepared thus far are both interesting and accessible to a broad audience and we are quite excited at the prospect that these might attract subsequent media attention for the researchers, for NCI and its stakeholders, and more generally for the role that high-end computing plays in contemporary, leading edge research.

Operations and Services

National Facility

A major focus of work for NCI over the past year has been the acquisition of the new peak facility. This commenced in March 2008 with high level discussions between ANU/NCI, BoM and CSIRO leading to a decision to proceed with a joint tender for an integrated computational environment for the national earth systems science community. The goal for NCI was a new peak research facility, the performance of which would be an order of magnitude (i.e., a factor of 10) greater than the current National Facility.

The tender was released in April 2008 and closed in late May 2008. Following this, the evaluation of tenders took place according to a detailed evaluation methodology developed by ANU, BoM and CSIRO staff. The evaluation was both complex and lengthy due to the need to marry the Bureau's requirements for an operational facility with NCI's more diverse, yet more spartan, needs for a high-end research system into an interoperable framework. This culminated in early October 2008 with the Joint Steering Committee (established with representation from BoM, ANU/NCI, CSIRO to oversee the procurement) recommending to each of BoM and ANU that the tender from Sun Microsystems be accepted.

In parallel with this, the NCI Steering Committee at its June 2008 meeting approved the acquisition of a small facility whose general architecture and performance characteristics would be comparable to that of the new peak system. This facility, intended to provide both a development platform for National Facility staff in preparation for the new peak system and a transition facility to enable researchers to develop, modify and test their codes in preparation for the transfer to the new facility, was ordered from SGI in late September 2008, installed in December 2008 and commissioned early in 2009. Since then, the SGI Altix XE has proven to be a very useful system, providing substantial performance increases for a diverse collection of codes and packages, and also serving as a platform on which National Facility staff could hone the software stack and the file system implementation of the new peak system.

Contract negotiations for the new peak system took place during the period from early November 2008 through until early March 2009 in a steadily deteriorating economic environment. The risk analysis undertaken in these circumstances was of considerable rigour and was substantially more robust than would have been required under more normal circumstances. In late February 2009, the Steering Committee, through a circulation of papers and email consultation, recommended to ANU that the goal of achieving a 10-fold increase in performance had to be delivered to keep faith with the primary project objective of providing Australian researchers with a system that would be internationally significant, despite the additional costs that would be incurred. With the completion of the risk analysis, ANU signed a contract with Sun Microsystems on 11 March 2009.

At the time of writing, the first phase of the new facility has just been commissioned on 16 September 2009, with a full service to the research community being available from 17 September. The former system (SGI Altix 3700) was shut down on 22 September and the new system, which is one-eight of its final configuration, is now shouldering the user load. In line with our expectations, the performance of the new system is excellent and initial users have commented very favourably on this, and are expectantly anticipating the commissioning of the full system at the end of the 2009. While Section 3.1.1 provides a fuller description of the procurement and installation process, it is important to note here that the new facility involves a number of leading edge developments including diskless compute nodes and a very high performance file system (based on the on the cutting edge of Lustre). That this has been achieved, particularly within a very limited timeframe is due to the collaboration between Sun and key staff of the NCI National Facility, the technical expertise of whom needs to be noted as a unique national resource and one which is held in very high regard within Sun Microsystems.

Specialised Facilities

The Specialised Facilities program was initiated with a call for expressions of interest in July 2008. At the close of the call in late August 2008, some seven applications had been received. Of these, the Evaluation Committee deemed two to be uncompetitive,

while two others were regarded as candidates for support under the Computational Tools and Techniques Program. The remaining three were then subjected to closer scrutiny, As is described in Sec 3.1.2, the Steering Committee ultimately approved the implementation of two Specialised Facilities, one in Bioinformatics, to be hosted by the University of Queensland, and another for the characterisation research community in Imaging and Visualisation to be hosted by Monash University. Regrettably, the implementation of the Specialised Facilities program has been delayed some months by an appeal against the evaluation process and the Steering Committee's decision concerning one of the facilities. Sec 3.1.2 details progress to date which is on track for a 2010 commencement of the program, with resources allocated through the NCI Merit Allocation Scheme.

Computational Tools and Techniques

The call for Expressions of Interest in the Specialised Facilities program during 2008 identified two areas, namely earth systems science (particularly climate science) and astronomy, which might be supported more appropriately under the CT&T program. In these circumstances, the Steering Committee decided to proceed with the implementation of the Specialised Facilities program and to subsequently implement the CT&T program. However, with contract negotiations for the peak facility proceeding in what was becoming an increasingly uncertain financial environment, it became necessary to delay the implementation of the CT&T programs in order to have some certainty concerning the cash flows needed for the peak system procurement.

At its 18 March 2009 meeting, the Steering Committee agreed to the implementation of the CT&T program with a budget of $2M for the period 2009-11.

As was described in the 2009-10 Business Plan, the program comprises support for:

- the earth systems science community, as previously identified in the Specialised Facilities Expressions of Interest, and consistent with the joint tendering framework for the new peak facility as well as the Super Science initiative announced in the 2009 Commonwealth budget,
- astronomy, which was also identified through the Specialised Facilities Expressions of Interest, and which is being developed in consultation with the research community through Astronomy Australia,
- data intensive computation, with an emphasis on providing these services through the emerging new cloud computing paradigm, and
- possible additional areas of research or research services.

Of these, the program in data intensive computing is already underway and good progress is being made.

Plans are also well advanced for the implementation of the climate science and astronomy support program from January 2010, consistent with cash flow considerations. At this time, project plans are being finalised and staffing agreements have been drafted and agreed to in principle by the host institutions. Sec. 3.1.3 provides greater detail about both the plans and progress towards their implementation.

3. PROJECT ACTIVITIES

3.1 Operations and Services

3.1.1 National Facility

Summary of Facilities

The NCI National Facility is hosted at the Australian National University under an agreement through which ANU provides computing services and specialist support for research users nationally. During the year covered by this report, the National Facility comprised a peak computing system, a SGI Altix 3700 ('ac') installed in 2005, and a cluster system. In the first part of the reporting period, the latter comprised the Dell Linux cluster ('lc') installed in 2003, but which was decommissioned late in 2008 to make way for the development / transition system, a SGI Altix XE ('xe'), installed in December 2008 and referred to in greater detail later in this section.

In addition, NCI has access to large-scale data management facilities through a Facilities Management Agreement with ANU. This agreement was updated and renegotiated in late 2008 and was approved by the Steering Committee at its November 2008 meeting..

Table 1: Summary of the two main systems available to users from January 2009.

SGI Altix 3700 BX2 (2005-09)		SGI Altix XE (Dec 2008-)	
Current Peak System		Development / Transition system	
CPUs: 1920 (Intel Itanium2 1.6 GHz) Main memory: 5.5 Tbytes Storage: 40 Tbytes Performance Peak: 14 TFlops Sustained: 21K SPECFP Resources: 16.8M hrs p.a. Purchase price: Approx $14M		Cores: 1248 (Intel Xeon 2.8 GHz) Main memory: 2.5 Tbytes Storage: 90 Tbytes Performance Peak: 14 TFlops Sustained: 12K SPECFP Resources: 10.6M hrs p.a. Purchase price: Approx $1M	

Usage of the Facilities: July 2008–June 2009

The use of the facilities is governed by a partnership model, the latest form of which was introduced at the beginning of 2009 following the signing of the NCI Partner Service Agreement by CSIRO and ANU. The reporting period FY2008-09 covers two rounds of merit allocation over the calendar years 2007-08 and 2008-09. Table 1 shows the share allocations over these two periods.

Table 2: NCI Share Allocations

Share	2007-08[1] Fraction	CPU/Core Hrs (Millions)	Share	2008-09[2] Fraction	CPU/Core Hrs
MAS	73.00%	12.764	MAS	45.99%	12.764
ANU	27.00%	4.720	ANU	22.72%	6.306
			CSIRO	22.39%	6.213
			CAWCR/CSIRO	3.79%	1.051
			Other partners	5.11%	1.417
Totals	100.00%	17.484	Totals	100.00%	27.751

Appendices 3 and 4 present the six monthly National Facility Operations Reports for the periods July–December 2008 and January–June 2009, with further information about the National Facility being available at: http://nf.nci.org.au/ .

In what follows, we focus on the Merit Allocation Scheme (MAS) which is provided for by the Commonwealth (NCRIS) funds.

[1] During the period July – December 08, the National Facility systems comprised the SGI Altix and a Dell Cluster.
[2] During the period January – June 2009, the National Facility systems comprised the SGI Altix 3700 and a SGI Altix XE.

Since a summary of MAS Allocations for the 2008 calendar year was presented in the previous annual report, the summary here is for the 2009 calendar year allocations only.

Implementation of the Merit Allocation Scheme for 2009

As was foreshadowed in the 2008 Progress Report, the MAS processes have been substantially strengthened this year to enhance the rigour and quality of the processes, given the magnitude and value of the resources being allocated. The revised processes were developed by the NCI Director, the MAS Chair and the National Facility Manager, and benefitted through consultation with the ARC and the Merit Allocation Committee (MAC). From 2009 onwards, the ARC has indicated its willingness to provide an observer to the MAC meetings to assist in the development of our processes and quality control.

The revised processes for the 2009 MAS round, described in the 2008 NCI Progress Report and also at http://nf.nci.org.au/accounts/#computationalprojects, were based on the following principles that were affirmed by the NCI Steering Committee during 2008:

- Research quality
  - *Research merit including the potential of the work to generate new knowledge in an important area, the comparative scientific merits of the work within its discipline, originality and international competitiveness.*
  - *National benefit and research priorities.*
  - *Experience and demonstrated research capacity of the applicant and the project team.*

- Appropriateness of the NCI computational resources
  - *The need for such resources to conduct the research.*
  - *Suitability of the system (hardware and software) and its operational environment to support the project.*
  - *Evidence or experience to demonstrate that the project will use the facilities efficiently.*

- Reasonableness of the level of resources requested
  - *Relative to the total amount available.*
  - *Needed to make adequate progress in the proposed research program.*

- Track record of the applicant in using the NCI computational resources where relevant
  - *In the case of proposals to continue a project, the record of achievement and efficient use of previous allocations.*

For this year's round, which was based on the revised processes and criteria, the Merit Allocation Committee[3] (MAC) scrutinised more closely the research for which NCI support had been requested, and also the track records of the researchers and their teams. These new processes were generally well received by the research community and the quality of applications, in general, was good. The additional information available was appreciated by members of the MAC which noted that the application framework was now more in line with the processes of the ARC. The additional data included information about the research contributions and experience of CIs, the research area classification data, as required in ARC applications, the relevance of the project to National Priorities and NCRIS capabilities, and other research support provided to the project. The form also requested that each CI list the previous five years of their publications that were relevant to the computational project, in order to highlight the research outputs. The research impact was scrutinised more closely in this round by considering citations for this published work. The Committee discussed a number of small improvements to the application form and assessment criteria. Superprojects were implemented more fully as a way in which clusters of projects could be managed under a single Lead Chief Investigator (CI). Superprojects provide a method for Lead CIs to partition time allocations for various components of the work and also provide the ability to promote early career researchers by allowing them to develop project applications under the umbrella of an established superproject.

Merit Allocation Scheme (2009) Outcomes

Demand in 2009 increased substantially over the previous year, rising from requests totalling 19.5M hours for 2009 to 28.2M hours for 2009. This represents a level of demand that exceeds supply by a factor of 2.2—up from 1.5 in the previous year. As in previous years, all applications seeking large allocations were read by all members of the Merit Allocation Committee, while all other allocations were read by a least two committee members. Applications were judged on how well they met the MAS criteria, together with their

[3] The membership of the Merit Allocation Committee is a matrix that incorporates geographical diversity and which encapsulates the range of disciplines / fields that are represented in the applications from the computational research community. For the first time, CSIRO had two representatives on the Committee. MAC members are distinguished researchers, all of whom have strong publication and grant records. Amongst the membership is an ARC Federation Fellow and two ARC Professorial Fellows. The present membership comprises Prof Brian Yates (Chair, Tasmania), Prof Geoff Bicknell (ANU), A/Prof Jim Denier (Adelaide), Prof Julian Gale (Curtin), Dr Tony Hirst (CSIRO, CMAR), Prof Alan Mark (Queensland), Prof Ross McPhedran (Sydney), A/Prof Louis Moresi (Monash), Dr Alf Uhlherr (CSIRO, IMT). The NCI Director (Prof Lindsay Botten) and the National Facility Manager (Dr Ben Evans) are advisors to the Committee, while Ms Judy Jenkinson of the NCI NF serves as Secretary to the MAC.

efficient and effective use of the facilities. Requests seeking increases of more than 10% over the previous year were more closely looked at. Once all applications had been assessed and allocated, the committee made a second pass through the allocations to ensure relativity and fairness.

The following table summarises the allocations by organisation.

Table 3: Summary of the NCI Merit Allocation Round for the 2009 Calendar Year

Organisation	Request 2009	Grant 2009	Grant/Request
University of Sydney	3,786,700	2,886,742	76.23%
ANU	2,803,000	1,609,993	57.44%
Monash University	2,920,660	1,517,697	51.96%
University of Queensland	2,894,798	1,333,993	46.08%
Royal Melbourne Institute of Technology	1,752,000	1,139,997	65.07%
University of Melbourne	2,079,000	1,104,000	53.10%
University of NSW	1,791,000	1,092,989	61.03%
University of Adelaide	1,626,700	1,081,197	66.47%
Curtin University of Technology	1,160,000	1,010,000	87.07%
CSIRO	2,507,000	679,000	27.08%
University of Technology, Sydney	703,300	390,000	55.45%
University of Tasmania	1,454,000	313,000	21.53%
University of Western Australia	411,000	250,996	61.07%
University of Newcastle	780,000	200,000	25.64%
La Trobe University	180,000	120,000	66.67%
Macquarie University	169,546	119,998	70.78%
Murdoch University	240,000	100,000	41.67%
ADFA	96,000	65,998	68.75%
Griffith University	81,000	55,997	69.13%
James Cook University	70,000	39,998	57.14%
Swinburne University of Technology	38,000	16,996	44.73%
University of Southern Queensland	40,000	16,000	40.00%
University of Wollongong	18,000	14,000	77.78%
University of South Australia	12,000	12,000	100.00%
Totals	27,613,704	15,170,591	54.94%

Table 4 reduces the data in order to summarise the requests and allocations by organisational community. The columns respectively show the proportions of requests and grants by community, and the ratio of these.

Table 4: Summary for of the 2009 MAS Allocations by Community

Summary (%)	Request	Grant	Ratio
Group of Eight Universities	66.67%	72.14%	1.08
Independent Research Universities of Australia	4.02%	2.74%	0.68
Australian Technology Network Universities	13.14%	16.82%	1.28
Other Universities	7.10%	3.82%	0.54
CSIRO	9.08%	4.48%	0.49
Totals	100.00%	100.00%	

The growing level of demand is shown in Figure 1 which characterises the evolution of growth in requests, grants and usage over a five year period spanning two generations of the peak facility (pre 2005: with the HP Alphaserver and post 2005 with the current peak system). The strong spike in demand for 2009 is evident and it is anticipated that strong growth will continue in future years, with the new peak system addressing this from 2010 onwards.

Figure 1: Evolution of grow in demand for the peak facility 2004-09

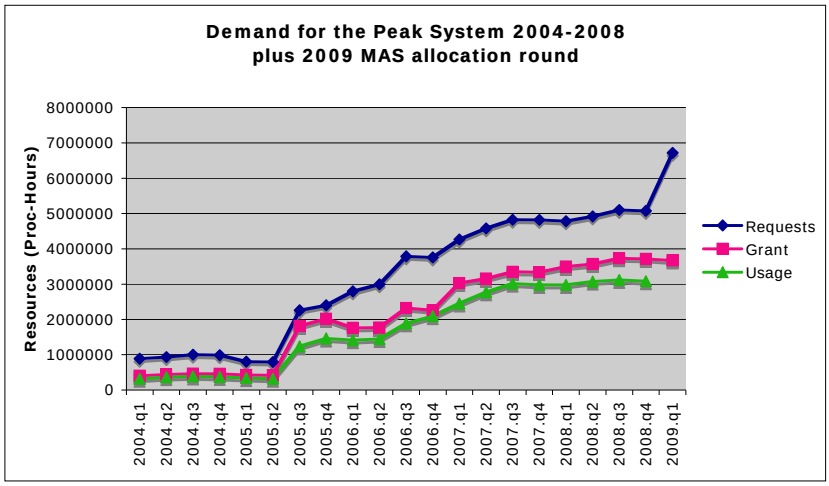

Figure 2: Resources allocation and actual usage by research fields for the MAS alone and all shares (MAS plus partners)

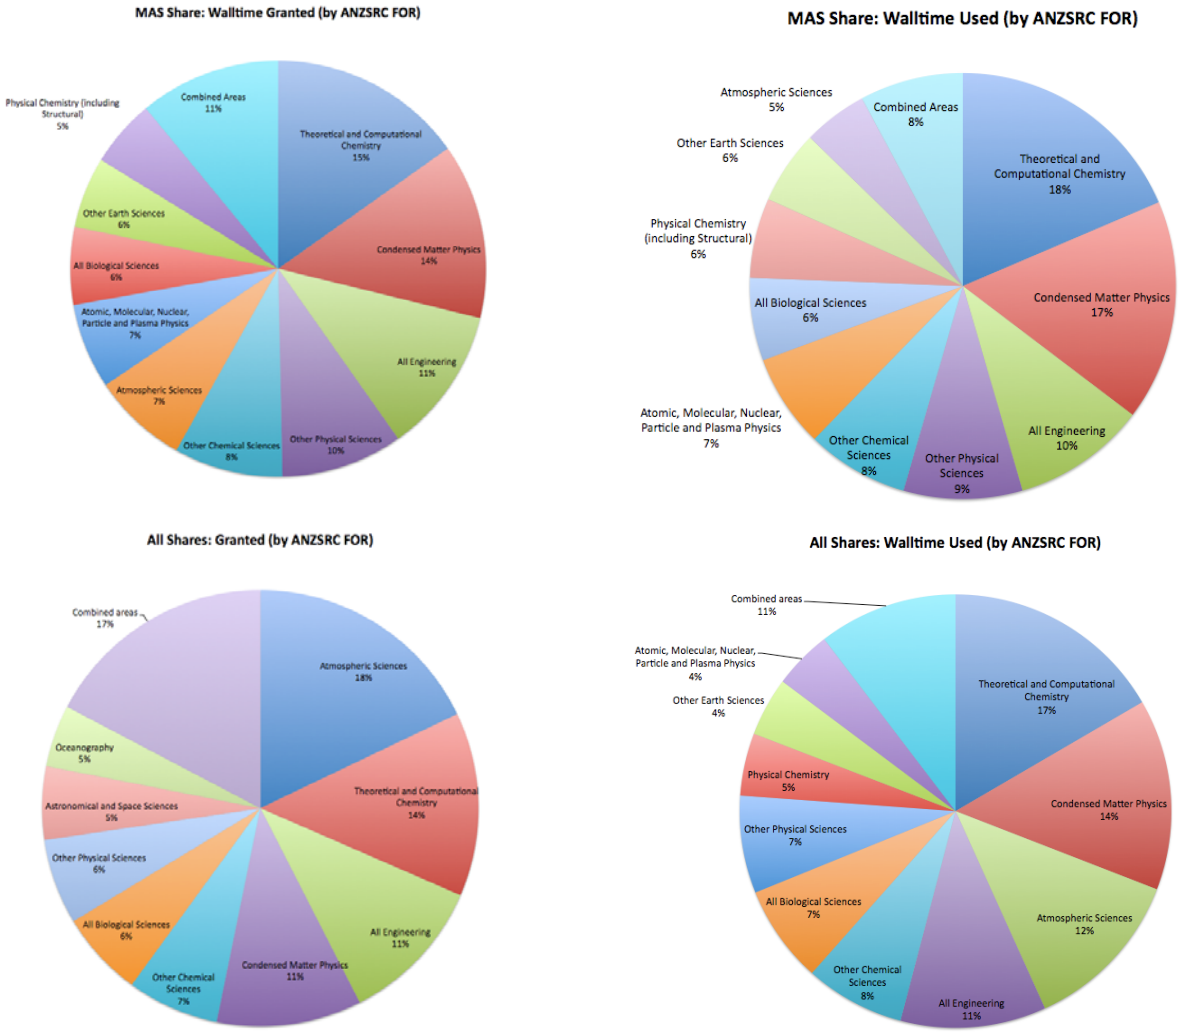

Figure 2 shows both the allocations for 2009 and total usage in 2009Q1 by research field. The two sets of figures display the outcomes for the MAS alone and the total allocations/usage by combining the MAS and all partner shares.

Summary Highlights of the 2009 MAS Allocations

- The largest grant for the 2009 calendar year was 920 kSU.
- A number of applications were not granted any time under the MAS.
- A total of 15.3 MSU was allocated by the MAC (representing an over-allocation of the MAS share of 20%).

  This compares with: 12.3 MSU allocated in 2007 (out of 19.2 MSU requested),
  and 13.9 MSU allocated in 2008 (out of 19.5 MSU requested).

- The amount allocated as a percentage of the overall MAS share is

	2006	2007	2008	2009
Top 10 projects	42%	40%	45%	42%
Top 20 projects	60%	58%	62%	60%
Top 40 projects	78%	80%	82%	80%

  In 2008 and 2009, project groups (superprojects) were considered rather than individual projects in the above table.

Table 5 Evolution of the Merit Allocation Scheme

Core Hrs (Millions)	2006	2007	2008	2009
Available	6.4	10.2	12.8	12.8
Requested	13.3	19.2	19.5	28.2
Allocated		12.3	13.9	15.3

Acquisition of the next-generation National Facility

Background

The procurement of the replacement of the National Facility peak system has been a prime focus of work for NCI during 2008-09.

High-level discussions in March 2008 involving ANU, BoM, CSIRO, and NCI led to a decision to proceed with a joint procurement for high-end computing facilities to provide for an integrated computational environment for the national earth systems science community—one that would span the needs of research, development and operational services. In this way, the Bureau of Meteorology would refresh its operational system, replacing its present NEC SX-6 facility based at its Melbourne headquarters, while ANU would update the NCI National Facility (NF) that serves as Australia's peak computational research system.

By choosing to proceed in this way, it was possible to achieve a higher level of functionality and collaboration than would have been the case by alternative means. The decision to proceed with a joint procurement also provided a framework for new investment in future years to address the urgent need for additional resources for climate modelling.

Development/Transition System

In parallel with the acquisition of the new peak facility, described immediately below, the Steering Committee recognised the need to:

- provide additional compute resources within the NCI partnership model prior to the commissioning of the new peak system (in late 2009), and
- establish a system with similar characteristics to that of the next generation peak facility, allowing the staff at the NCI NF an opportunity to build expertise in configuring, managing and operating a large scale production cluster with an Infiniband interconnect and a Lustre filesystem.

At its April 2008 meeting, the NCI/SC approved the acquisition of a development/transition system, allocating $800K to this purchase, as had been foreshadowed in the 2007-08 Business Plan. Subsequently, a request for tender for this system was released on 21 June 2008, with tenders closing on 17 July 2008. During this process, CSIRO/CAWCR indicated that it wished to acquire additional dedicated resources for climate modelling in the lead up to the IPCC 5th Assessment Round and proposed an additional injection of $150K from CSIRO, bringing the total funds available for this procurement to $950K. Following the evaluation of tenders (from six vendors), a SGI Altix XE system was selected, with the Steering Committee agreeing to its acquisition in September 2008. A contract was let in late September 2008 and the new system was installed in November 2009 and commissioned in January 2009.

The XE system, the specifications of which are listed earlier in the document, has been available to the user community since February 2009. Since then, it has proven to be a fine system yielding significant performance gains over the older SGI Altix 3700 system on a range of codes. It has also provided a platform for National Facility staff to develop the software stack to be installed on the new peak system and to acquire substantial experience in the management of a Lustre-based filesystem, as well as providing a facility by which users can port and test their codes in an environment which is almost identical to that of the new peak facility. As a consequence of

this, the transition process from the old to the new peak systems has been able to be accelerated substantially and has been largely transparent to the user community.

Procurement Process for the Peak Facility

The joint procurement of interoperable peak facilities by the Bureau of Meteorology (BoM) and ANU took place by a tender let through AusTender and led by BoM. Despite the orientation of the tender towards earth systems science, arrangements were put in place through the selection criteria to ensure that the NCI National Facility system would deliver strong benefits to the broad span of research areas supported under the NCI Funding Agreement.

The tender was released on 3 April 2008 and closed on 29 May 2008, after which a comprehensive evaluation process, conducted under a closely documented Evaluation Plan and Probity Plan took place in ensuing months. This process was overseen by a Joint Steering Committee comprising equal representation from BoM and ANU/CSIRO/NCI constituencies, while the evaluation was undertaken by an Evaluation Team led by Mr Tim Pugh (BoM) and Dr Ben Evans (ANU/NCI) with representation from ANU, BoM and CSIRO.

The decision to proceed with the tender from Sun Microsystems was finalised in early October 2008, with the Joint Steering Committee formally communicating its recommendation to each of ANU and BoM. From then through until early March 2009, BoM and ANU pursued separate contract negotiations with the vendor, although maintaining the vision of interoperable facilities through ongoing communication, implicit in the framework of the Joint Steering Committee that had been adopted. These negotiations were both lengthy and detailed, given the complexity of a joint procurement and also the growing uncertainty within the global economic environment. Indeed, contract negotiations took place in a steadily deteriorating and increasingly volatile economic environment in which the Australian dollar dropped in value from approximately US$83c in October 2008 to around US$62c in early March 2009. As a consequence of this, the cost of acquiring a system that met the goal of an order of magnitude performance enhancement increased by approximately 20-25% during that period. Furthermore, the risk analysis and due diligence processes that were needed were significantly more rigorous and robust than would have been required in more conventional circumstances

The NCI Steering Committee was consulted out of session (February 2009) through the circulation of a paper that sought their agreement for an increased level of expenditure (above that budgeted for in the 2008 Business Plan) in order to meet the original performance goal. The paper apprised Committee members of the risk analysis that had been undertaken, with this including realistic and worst case projections of partner income and the examination of mitigating strategies such as the rescheduling or reduction in the scope of subordinate programs (i.e., Specialised Facilities and CT&T), the possible invocation of a change order with the vendor (to reduce the scale and cost of the acquisition), and the removal of uncertainty by hedging exposure to currency fluctuations through forward buying of US currency.

Table 6: Comparison of the current and next generation peak systems

SGI Altix 3700 Bx2 (2005-09)		Sun Constellation System (2009-)
CPUs: 1920 (Intel Itanium2 1.6 GHz) Main memory: 5.5 Tbytes Storage: 40 Tbytes Performance Peak: 14 TFlops Sustained: 21K SPECFP Resources: 16.8M hrs p.a. Power: approx 300 kW Purchase price: Approx $14M		 	Cores: 11936 (Intel Xeon 2.93 GHz – Nehalem series) Main memory: 36 Tbytes 7x Storage: 500 Tbytes 12x Performance Peak: 140 TFlops 10x Sustained: 240K SPECFP 12x Resources: 110M hrs p.a. 6.5x Power: 604 kW 2x Purchase price: Approx $14.5M

With all Steering Committee members recommending the increased allocation of funds to ensure the performance objective for the system was met, NCI advised the Department of Innovation of its plans and briefed government of the attendant risks. During that meeting, the Department offered to accelerate the payment of the 2009 tranche of funds ($9.5M), by bringing forward the payment date from November 2009 to July 2009. This was subsequently formalised in a letter from Ms Anne-Marie Lansdown to the NCI Chair (21 February 2009), and accepted in a letter of reply (5 March 2009), noting that this would be of significant assistance to NCI as it would reduce the costs involved in the forward acquisition of US currency and assist in the management of cash flows. In keeping with the Government's requirements, appropriate new milestones were subsequently agreed with the Department in May 2009.

It was not until March 2009 that both ANU and BoM were able to complete negotiations and to sign their respective contracts. In the case of the NCI Facility, ANU signed its contract with Sun Microsystems on 11 March 2009. Immediately thereafter, ANU confirmed NCI's exposure to currency variations and made arrangements to acquire US currency to meet the future schedule of payments.

A public announcement of the signing of the contracts was made in two press releases (one from ANU and BoM, and another from Sun Microsystems, also involving ANU and BoM) on 18 March 2009.

Site Upgrades

Taking place in parallel with the contract negotiations was substantial background work in preparing for the necessary power and cooling upgrades in the NCI National Facility Data Centre located in the ANU Leonard Huxley Building. Due to the long lead times in the production and shipping of major water cooling plant from overseas suppliers, and also due to the length of the contractual negotiations with Sun, the tolerances on delivery schedules for power and chiller equipment became quite fine. The plant upgrades total approximately $4.5M and are able to provide for the facility listed in the contract with Sun, plus possible future upgrades of up to 50% of the planned capability.

Installation and Commissioning of the New Sun Constellation System

Due to time, space, power and cooling constraints, the installation and commissioning of the new peak facility is occurring in two phases.

The first phase installs one-eighth of the final facility and is physically quite small, requiring only five standard racks. It will operate within the power and cooling envelope of the current machine room, and has been installed in parallel with preparatory site work (i.e., new power and cooling infrastructure for the main system) in a new purpose-built building in a former car park behind the Leonard Huxley building. The first phase allows users to be migrated from the old SGI Altix 3700 system, after which it can be decommissioned, the space that it occupies recovered, and critical work in the machine room (i.e., upgraded power and cooling infrastructure) to take place before the bulk of the new facility is installed.

Some delays in the installation schedule arose as a consequence of Sun's inability to provide solid state disk drives in the metadata servers of the filesystem until 2010. Understanding the implications of this, and determining a resolution via a Change Order (which involves a contract variation) took a number of weeks and consequently delayed production and shipping schedules.

Phase1 of the Installation

The hardware delivery for Phase 1 of the installation took place over a one week period commencing on 24 July 2009. By mid-August, the system comprising five racks (2 C48 processor racks, 1 M9 Infiniband rack, and 2 racks of disk storage) had been assembled and cabled, and had passed a suite of hardware diagnostic tests. National Facility staff was provided with access to the system from mid-August in order to assemble and test the software stack, to configure the Lustre filesystem, and to test a range of user codes. The final stability and acceptance test of 120 hours of continuous running commenced at 9.00 am on Wednesday, 9 September and was completed successfully at 9.00 am on Monday, 14 September. ANU was provided with Acceptance Certificates on Monday, 14 September and the system was formally accepted by ANU on Wednesday, 16 September 2009, triggering the first major payment to Sun for the new facility.

Although the new system comprises only one-eighth of the final configuration, it has a sustained throughput that is 50% greater than the entire SGI Altix 3700 which it replaces. The first service to users of the facility was provided from 17 September 2009, with all users being transferred by Friday, 18 September. The SGI Altix 3700 was switched off and decommissioned on Tuesday, 22 September 2009.

The overlap between the commissioning of the new system and the decommissioning of the old facility has been necessarily brief due to delays that have occurred previously. Any longer transition period would have led to downstream delays in the planned refurbishment of the data centre and the subsequent installation of the remainder of the new Sun facility, thus delaying the commissioning of the new system until early in 2010.

Despite the brevity of the transition period, there has been no apparent inconvenience to the user community. With the new SGI Altix XE (i.e., the development / transition system available from January 2009) being essentially identical in architecture to the new Sun Constellation system, and having been available to the user community for almost the whole of 2009, there has been sufficient time for NCI users to test and redevelop codes in this new environment.

The following comments received from researchers using the new system (vayu.nci.org.au) are indicative of the transparency of the transition process and of the performance of the new system.

> "Thank you for such a smooth transition to VAYU. ... our ... code runs three times faster on the VAYU"
>
> "So far I've built our software without any difficulty. The whole operating environment is indeed practically

identical to the (development) machine so the transition to the new system appears to be trivial."

"The results from the new system are quite spectacular"

"This is very impressive, and beats any comparable system by a fair margin."

"Thank you very much for the all information. It certainly sounds like you folks have plenty of experience ... You can be sure I'll be in touch :-)"

"I couldn't resist putting the CSIRO ... model through its paces on vayu as soon as it became available. The results are very impressive"

"So far, so (very) good... I'm seriously impressed: from my perspective at least, it would not have been possible for you to make the transition to vayu any smoother or easier. Thank you!"

The delays referred to above have caused some minor additional expenditure due to the need to extend the maintenance on the SGI system and its cooling plant by a fortnight. The approximate cost of $15,000 is well within the contingency built into the budget.

Phase 2 of the Installation
Following the decommissioning of the SGI Altix 3700 on 22 September, the machine has been removed in preparation for the refurbishment of the computer room. This includes the removal of the old computer room floor, the removal of old plumbing (associated with the SGI Altix 3700 water cooling), the installation of new underfloor plumbing for the higher capacity cooling required by the new Sun system, and the installation of a new computer room floor with a higher load bearing capacity. This work should be completed by mid-October.

The old SGI, which is a shared memory system, has no future value to any prospective purchaser (unlike the previous HP/Compaq cluster system which could be sold in discrete units). Accordingly, it is being disposed of for scrap value only, with no net cost for its removal accruing to ANU/NCI.

The delivery of Phase 2 hardware, which comprises the bulk (7/8th) of the new facility, is anticipated to commence in the week commencing 19 October 2009. That week will see the delivery of the filesystem hardware, with the 14 C48 compute racks arriving in the following week, and the cooling hardware (i.e., the rack doors) arriving in the week after that. A phased delivery of this kind is necessary because of the limited space for storage. In parallel with this, the new power and cooling plant (in a new purpose built building) will have been commissioned. By mid-November, the Phase 2 system will have been assembled and cabled, in preparation for system diagnostics which are scheduled to commence on 12 November. Acceptance and stability testing is scheduled to commence in mid-December and to be completed before Christmas 2009, at which time the new system will be available to the research community.

At that time, users will be migrated from the smaller Phase 1 system onto the larger Phase 2 system, after which the Phase 1 system will be shutdown and physically moved and integrated with Phase 2 to form the final merged system. This integration is anticipated either late in December or early in January 2010 and can be undertaken with disruption to the service provided by the Phase 2 system

3.1.2 Specialised Facilities

The Specialised Facilities Program provides a mechanism by which NCI can extend the range of computational services available to the Australian research community by investing in facilities operated by other specialist providers of high-end computational resources.

This is particularly relevant when either:

- the NCI National Facility is not necessarily well suited (or efficiently suited) to the task (e.g., the service does not require the high levels of connectedness of the NCI NF system), or
- it would not be an effective use of resources to provide the service (software and support) through the NCI National Facility, if this involves duplicating resources and services already available elsewhere in Australia.

To meet the need for such requirements, NCI has allocated $2.4M to be used for acquiring shares in computational services available through other specialist providers. The intention of this is to invest in two specialised facilities, and to make these services available through the NCI Merit Allocation Scheme.

A broad call for Expressions of Interest was issued on 21 July 2008, with some seven applications received by the closing date at the end of August 2009. These were assessed according to the following selection criteria:

- the extent of the demand for the facility by 'high-end' users in the Australian research community, including the NCRIS capability areas,

- the access to the specialised facility and user support that will be provided for researchers Australia-wide,
- the characteristics of the specialised facility and the manner in which it complements the capabilities of the National Facility,
- the proposed investment in the specialised facility represents good value in comparison with other possible investments,
- the host organisation is capable of providing the required level of service to users as a component of the 'national advanced computing infrastructure',
- the host organisation is able to implement policies, system scheduling and processes to accommodate projects granted resources under the NCI Merit Allocation Scheme.

Implicit in these criteria was that funding provided under the Specialised Facilities program would be tied to meeting the following:

- acceptability to the Department of Innovation,
- an appropriate level of co-funding,
- put in place through contractual agreements linking accountability to NCI.
- arrangements for allocation of associated resources through the NCI MAC.

A two stage selection process was adopted, in the first phase of which uncompetitive applications were culled and applications that could be supported in other ways identified. In this way, a shortlist of three competitive applications was selected for closer scrutiny, and a report provided to the Steering Committee. Of the seven applications, three were deemed uncompetitive, and two, one from each of climate science and astronomy, were regarded as candidates for support under the Computational Tools and Techniques program. Three applications, two from Bioinformatics and one in Characterisation and Visualisation, were taken forward for closer analysis.

The second phase undertook a closer examination of the shortlisted applications, with a number of clarifications sought from, and provided by, their proponents. The NCI Steering Committee was then provided with a final report recommending that one of the two applications from Bioinformatics be progressed to funding, and that further investigation of the demand for the Imaging and Characterisation Facility be undertaken before finalisation of a funding recommendation.

At its 28 November 2008 meeting, the NCI SC passed both recommendations and the Director subsequently communicated the outcomes to proponents.

An appeal against the decision of the Bioinformatics facility, however, delayed the implementation of the program at that point. The Steering Committee Chair determined that to ensure that all parties had confidence in the decision and the decision making processes, a formal review should be established. This was set up in early February 2009 with Processor Ian Sloan, AO, FAA (Scientia Professor of Mathematics at UNSW) being requested to undertake a review of the selection processes. Professor Sloan, after reviewing all relevant documents, provided a report to the Steering Committee Chair in early March 2009 recommending that the appeal be rejected and that the original decision be upheld. His report was considered at the 24 March 2009 meeting of the Steering Committee which subsequently recommended that the Director proceed with the implementation of a Specialised Facility in Bioinformatics at the University of Queensland.

In parallel with this, NCI continued to work with the proponents of the Specialised Facility in Imaging and Visualisation to finalise a surveys of needs and demands for these services. This survey was undertaken by the Director, Mr Rob Cook (as a consultant to Monash University) and Mr Richard Farnsworth (Australian Synchrotron), and involved meetings and teleconference discussions with the Australian Microscopy & Microanalysis Research Facility (AMMRF), the National Imaging Facility and ANSTO. At the conclusion of these meetings, a report of the Needs Survey was prepared by Mr Cook and was considered subsequently by the original Specialised Facilities Evaluation Committee. This committee convened by email exchange and agreed that the demand had been sufficiently well articulated to proceed with implementation. This recommendation was considered and agreed to at the March 2009 meeting of the Steering Committee which requested that the Director proceed to work with Monash University on its implementation.

Following this, a services agreement was drafted by the Director and the ANU Legal Office in April-May 2009. This document takes the form of a sub-contract to the NCI Funding Agreement and is consistent with sub-contracts under NCRIS in other capability areas. It is also consistent with the NCI Funding Agreement, ensuring that particular clauses in the original agreement flow through into the sub-contract. The agreement has been structured in such a way that the body of the two (Bioinformatics and Imaging and Visualisation) contracts are identical, with differences confined to the Schedules which define the facilities and services to be provided and offered through the NCI Merit Allocation Scheme.

On the basis of progress reached in developing the agreements for Specialised Facilities program (together with associated progress in the development of the CT&T program), the Commonwealth has paid the full tranche of 2009 funds some months early. Implicit in this payment, on the basis of the agreed milestones, is that the Commonwealth has implicitly agreed to the sub-contracting arrangements (as referred to in the NCI funding agreement).

At the time of writing, the contract negotiations are almost complete. The bodies of both contracts are largely agreed and we are awaiting confirmation about the final batch of changes. The only significant change that has arisen has been a request from the

University of Queensland that QCIF (Queensland Cyber-Infrastructure Foundation), which is the channel though which Queensland Government coinvestment flows to this initiative, also becomes a signatory to the contract. NCI understands from UQ and QCIF that unless this is the case, Queensland Government funds cannot be allocated to this initiative. Accordingly, at the request of the primary service provider (UQ), the contract has been modified to accommodate two service providers (UQ and QCIF). Apart from this modification, the contracts for the two specialised facilities are identical.

Table 7: Summary of the arrangement for Specialised Facilities

Specialised Facility	Lead Organisation(s) / Signatories	Anticipated Coinvesting Partners (cash / in-kind)	NCI Investment / Period of Investment	Service Length / Anticipated 2010 Shares
Specialised Facility in Bioinformatics	University of Queensland / UQ and QCIF	UQ, QCIF, QFAB, CSIRO, NCI	$1,200,000 2 years (2010-11)	3 years (2010-12) 21.2%[4]
Specialised Facility in Imaging and Visualisation	Monash University	Monash, Australian Synchrotron, VPAC, CSIRO, NCI	$1,200,000 2 years (2010-11)	3 years (2010-12) 14.3 %[5]

The Schedules are also close to completion, although there is still some refinement required in the definition of the resource unit[6] in both contracts. In the case of the Monash agreement, there is also some additional information required about the range of services to be provided, and particularly about the split between online and batch usage in visualisation and image analysis, and the manner in which researchers would access and use these two modes. It is expected that agreements will be finalised during October 2009, in preparation for the commencement of their services in 2010 in accordance with cash flows.

In either case, there are very strong levels of coinvestment, although these are yet to be finalised. NCI's investment in both facilities is $1.2M. This will be paid over a two year period (within the period of the NCI Funding Agreement), although the services will run for a period of three years. CSIRO is investing an identical amount over an identical period in each of two facilities.

3.1.3 Computational Tools and Techniques Program

The Computational Tools and Techniques (CT&T) program is intended to support the development of software tools to improve particular applications or to provide targeted user support for particular user communities. As with the Specialised Facilities Program, the primary driver for the establishment of a CT&T project is the need and demand from particular research communities.

The call for Expressions of Interest in the Specialised Facilities program during 2008 identified two areas, namely earth systems science (climate science) and astronomy, which might be supported appropriately under the CT&T program. Accordingly, it was decided to proceed with the implementation of the Specialised Facilities program and subsequently to implement the CT&T program. At that time, however, with contract negotiations for the peak facility taking place in an uncertain economic environment, it became necessary to delay the implementation of the CT&T program in order to have certainty about the required cash flows.

At its 18 March 2009 meeting, the Steering Committee agreed to the implementation of the CT&T program with a budget of $2M for 2009-10. The program comprises support for:

- the earth systems science (previously identified in the SF EoI) and consistent with the joint tendering framework for the new peak facility, as well as the Super Science initiative announced in the 2009 Commonwealth Budget,
- astronomy (also previously identified in the SF EoI), to be developed in consultation with the research community through Astronomy Australia,
- data intensive computation, with an emphasis on implementing this through cloud computing, and
- possible additional areas of research or research services.

At this time, the activity in data intensive computation / cloud computing is well underway, with solid progress also made in preparing for the implementation of the activities in supporting the earth systems science, astronomy and environment communities. Cash flow considerations require that the bulk of the activities occur in 2010-11. This timeframe, however, is particularly appropriate for both the earth systems science and astronomy communities in the lead up to new facilities becoming available in future years under the Super Science initiative.

[4] Anticipated as at the time of writing, but requires confirmation of other stakeholder inputs.
[5] Anticipated as at the time of writing, but requires confirmation of other stakeholder inputs.
[6] The resource unit seeks to define equivalence between processor time (traditionally used as the resource metric) and data storage, which is becoming increasingly significant as research becomes more data-intensive. NCI, Monash and UQ are working to develop a broadly acceptable metric that takes into account the range of storage requirements that are anticipated.

Progress in each of these areas is summarised under the relevant heading below.

Data Intensive Computation / Cloud Computing

This activity has a total allocation of $600K, comprising 1.5 EFT staff positions in the National Facility and up to $75K p.a. in associated support costs (e.g., funds to acquire resources on commercial cloud systems), for a period of two years.

Work commenced in July 2009, with two experienced National Facility staff—Joseph Antony (0.5 EFT), and Ahmed El Zein (1.0 EFT) assigned to this program. Since then substantial preliminary work has been undertaken, particularly the implementation of virtualisation tools on the National Facility Data Cloud (i.e., the system dc.nci.org.au) to provide a platform for trialling cloud software for use in a data intensive environment. In particular, the NF team is trialling the Eucalyptus suite which is an open-source system for implementing on-premise private and hybrid clouds using the hardware and software infrastructure that is in place. It provides functionality similar to that available on the commercial Amazon Elastic Cloud (EC2). The group has also been working towards the implementation of Hadoop and MapReduce which is a suite of open source tools for the development of scalable, distributed computing, particularly for data intensive computation.

The team has been also implementing tools to handle the transfer of data from external stores (e.g., the BoM/CSIRO store of climate data in Melbourne) and instruments (e.g., SkyMapper in optical astronomy). The group is also preparing for the Australian implementation of a climate data (PCMDI) super-node.

Earth System Science

At the March 2009 meeting of the Steering Committee, members agreed to an allocation of $300K p.a. (i.e., 2 EFT) to provide support to the ESS research community for a period of two years.

A draft project plan has since been developed in consultation with the Earth Systems Science (ESS) community, focussing on two key areas of support for the ACCESS project—specifically the extension, maintenance and continuous upgrading of the ACCESS model on the NCI National Facility, and the development of stronger ties between the university climate community, CAWCR and the Hadley Centre in the UK.

Following consultation, there is community-wide agreement that one of the two positions will be based in the National Facility and that the second will be based in Melbourne, the hub of the development of the ACCESS suite. This latter position will be hosted by Monash University, with the incumbent to be supervised by the leading climate scientist, Professor Christian Jakob. A staff agreement has been developed by NCI and has been agreed to in principle by Monash University.

At this time, NCI is refining the goals of the Project Plan and will have the necessary agreements in place in October 2009, to allow for appointments to be made and a formal start to the project in January 2010. This project dovetails closely with the development of the investment plan for the new climate facility (Super Science initiative) and will also provide valuable support to a likely future application from the university climate community for an ARC Centre of Excellence.

Astronomy

The NCI Director has been working with the HPC Working Group of Astronomy Australia to develop a project plan for support for the Astronomy community.

Through this project, NCI seeks to seed substantial computational advances for the astronomy community, either enhancing the research capability of existing users of the computational facilities (and elevating their research expectations into the international rank), or springboarding new users from the astronomy community into the use of advanced computational techniques. Accordingly, the NCI Director and Manager of the NCI National Facility are working with a group led by Prof Matthew Bailes and Dr Darren Croton of Swinburne University, and are iterating a number of paradigm projects that will naturally foster both new users and new uses of the national HPC facilities, and which also will have a long term outcome in terms of enhanced levels of usage and sophistication of high-end computing in astronomy.

As is the case with the ESS CT&T project, there will be two positions funded in each of the two calendar years 2010-11. One of the positions will be located in the National Facility and the other at Swinburne University. The discussions at this time suggest that the Canberra based position will have a focus on optical astronomy applications and modelling, while the Swinburne position will be oriented towards radio astronomy and its modelling. Also being considered are applications of GPU computing which the astronomy community is keen to trial.

3.2 Planning, Access and Policy

3.2.1 Planning and Policy

Commercial Access

Consistent with NCRIS guidelines and past practices in APAC, the antecedent of NCI, the NCI Steering Committee has considered a policy for providing commercial access to NCI facilities as both a mechanism for fostering relationships with industry, and also as a source of external income for developing NCI resources and services.

Commercial access will be facilitated through a dedicated NCI Share of the resources (a small share of the total available resources) to provide the headroom needed to accommodate both:

- commercial opportunities, and
- access by new research teams, new applications or special project requirements, not previously granted resources under Merit Allocation.

If not used, this share would be available for redistribution to increase pro-rata the resources available to the Merit Allocation Scheme and Partners.

With resources on the NCI National Facility provided for from Commonwealth (NCRIS) funds and through co-investment by partner organisations, the Steering Committee has considered the possible interplay between partner usage of the facilities and commercial activities and has put in place a code of conduct based on a number of principles.

Underlying the policy on commercial access is the fact that partner engagement in NCI is predicated on the collaborative development and use of high-end services and resources through cost sharing in a way that leverages stronger outcomes than could be otherwise achieved.

Accordingly, the Steering Committee has affirmed that partner shares should be used in a way that is consistent with both the NCI Funding Agreement and the NCI Partner Service Agreement, with the latter explicitly referring to:

- keeping NCI informed about new opportunities in high-end computing that could be undertaken by or through NCI,
- doing all things necessary to allow NCI to perform its role in an effective manner.

In summary, the key principles guiding commercial access, potentially through partner engagement, are as follows:

- Commercial activities which generate a profit should be undertaken through the NCI share, except in limited cases in which there may be reasonable strategic reasons not to do so.
- Partner shares should not entitle partners to act as intermediaries that generate profits solely for their particular benefit since this would be contrary to (a) the spirit in which funding has been provided by the Commonwealth (under the NCI Funding Agreement), (b) the partner leveraging of Commonwealth funds that is explicit in the NCI Resource Allocation Model, and (c) the tenor of the NCI Partner Service Agreement, referred to previously.
- The provision of commercial access requires sufficient resources to be available so that any such service does not impact on other usage, and particularly the research usage of the facilities. Normally, such access would be provided from the NCI Share. However, in the event that sufficient resources are not available, new facilities may be acquired, either by:
    - NCI, provided that a contract is agreed between ANU (representing NCI) and the commercial client, ensuring that any NCI outlays are matched by an appropriate income stream, or
    - Partners, which may increase their partner share to accommodate commercial opportunities (consistent with the reasonable strategic relationships, outlined above) by either increasing their investment in NCI to acquire access to additional resources if they are available, or by providing additional funds to NCI to provision sufficient resources to be dedicated to the commercial task. In the case of the latter, the partner's share would be incremented by the annualised value of the additional resources provided at the conclusion of the commercial task.

In addition, the Steering Committee has agreed that any commercial usage of NCI facilities:

- must be consistent with the high-end services mission of NCI, necessitating the use of high-end infrastructure and services,
- must not breach licence conditions for the use of software packages on NCI systems (with this possibly requiring the procurement of appropriate licences by the client, or by NCI on behalf of the client),
- must not breach AARNet network access policies, particularly in relation to the transmission of commercial data to off-net sites,
- should be charged for at the market rate (set by NCI) reflecting full cost recovery including depreciation and consistent with the NCRIS Access and Pricing Principles.

The pricing of any commercial access will be implemented in a transparent manner that reflects the Competitive Neutrality guidelines that require that a competitive charge is levied—one that neutralises advantages associated with the tax-free status of the Host Organisation and shared infrastructure and resources funding which typically underwrite indirect project costs. This accords with the NCRIS Access and Pricing Guidelines which require that charges should reflect full-cost recovery.

3.2.2 Access

Access to NCI resources is provided in two ways:

- Merit and Priority Access which is funded by Commonwealth resources provided through NCRIS,
- Partner Allocations, in which NCI partners and affiliates acquire resources (respectively a defined share or a fixed quantity of resources).

These are dealt with under separate headings below.

Merit and Priority Allocation

Background

Merit principles lie at the heart of the operations of NCI. Over a number of years, these have evolved and been refined, but throughout they have comprised the key elements of research quality, the appropriateness of the computational resources for the task, the need for the requested resources to yield significant research progress, and the capacity of the applicant(s) to use the requested resources effectively and efficiently.

In a research environment of increasing global collaboration and competitiveness, NCI recognises the need to make changes to the Merit Allocation Scheme, bifurcating it into a General Access Scheme and a Flagship Program to ensure that the mission of NCI, namely *to provide Australian researchers with world-class high-end computing services* is met.

These changes were foreshadowed in the Business Plan and were driven by the need to better provide for the needs of cutting edge research. Since then NCI has been undertaking policy development work to develop the latter, the role of which will increase in significance if Australian research is to make best use of the rapidly increasing access to high-end computing services.

In this setting, the proposed Flagship scheme must deal with priorities set by research communities, research institutions and by government funding agencies. In this environment, the current Merit Allocation Scheme would become the General Access Scheme underpinning the breadth of computational research support. Both the General Access and Flagship schemes would be underpinned by the same well regarded merit principles that have been evolved over the years.

Goals of a Flagship Scheme

The overall objective of the Flagship Merit Allocation program is to increase the global competitiveness of Australian computational research, in some cases by providing a platform through which research teams can participate in the arena of grand challenge problems. For this to be possible, researchers need access to resource allocations which are comparable with those provided through the US Teragrid program. However, there is little in Australian computational research at present that is of a level of sophistication comparable to that seen in the USA or Europe. For example, at present only 6% of the National Facility exploits one-eight of more of the system while, in contrast, the usage patterns of major US systems (which are substantially larger than the new Sun system presently being installed) see a pattern of usage in which approximately half of the system time may involve the use of half of the system or more.

For Australian computationally based research to reach this scale of usage which exemplifies the cutting edge of international research, the availability of resources is not alone sufficient. In parallel with this, there needs to be a substantial upscaling of computational expertise to support Australia's leading projects. Such expertise is in short supply and needs to be used in a highly targeted way if the new resources coming on stream are to be used most effectively.

While NCI can target its expertise to support research projects, the targeting of the research to be supported, however, is not directly the province of NCI. Instead, decisions of this nature need to emerge from research communities and the research institutions, also informed by the strategic directions of government funding agencies. In such a process, NCI has an important role to play in applying merit principles around the determined priorities and assessing the computational suitability and merit of the tasks. The task of setting priorities and ensuring input from leading communities and institutions is seen as the role of a Science Advisory Committee, the importance of which was identified in the 2009-10 NCI Business Plan.

Overall, the specific goals of the Flagship program, within a priority framework set out as described above, would be:

- to substantially enhance the outcomes and international impact of leading Australian research through high-end computing,
- to upscale the sophistication of Australian usage of peak computational systems, particularly the use of parallel computation,
- to elevate Australian capability computing, and its use, into world class.

Intended recipients of Flagship Grants

The intended applicants for, and recipients of, Flagship MAS grants are:

- projects from single Chief Investigators or small teams undertaking high profile, internationally competitive research,
- substantial mission oriented collaborations such as Centres of Excellence, or their equivalents,

- research communities (e.g., climate science, astronomy, etc), which perhaps are members of international collaborations (e.g., using common codes) and whose outcomes are focussed on particular goals or national objectives.

In all cases, applicants would need to make the case for access to high terascale, and ultimately petascale, facilities.

Features of the Flagship MAS

At this formative stage, it is difficult to be certain about the features and the mix of services required in a Flagship MAS. Nevertheless, an analysis of what is required to move up "several notches" towards world's best practice (exemplified by the Teragrid, amongst others) suggests that the Flagship program should accommodate the following features / requests:

- Size and nature of the request (for both computational and data, treated in a holistic way),
- Multiyear allocations,
- Higher priority of access (to meet particular research goals / outcomes),
- Access to advanced support (needed to upscale the use of the facilities, e.g. parallelism),
- New services (to be suggested by the research community).

Criteria for Eligibility

In keeping with its intended standing as the pre-eminent merit allocation program, the Flagship MAS would exhibit the following characteristics.

- Demonstrated science excellence and research standing
  - Applicants would be required to demonstrate a portfolio of excellence measures which certify the quality and international competitiveness of the research, and of the chief investigators in the team. Mission-oriented collaborations and communities would need to provide corresponding evidence, commensurate with their size.
  - The use of pre-existing excellence measures such as a funding portfolio, or an association with a national research program, together with strong publication profiles, reduce the need for close (external) refereeing of the research/science program (as adopted by the Teragrid), allowing the NCI MAS refereeing to focus on the computational aspects of the application.
- A persuasive case for support under the Flagship Program, i.e.,
  - superior research questions / outcomes, higher impact results, deeper research questions posed / answered,
  - enhanced sophistication of the use of computational processes (leading to superior outcomes), e.g., higher levels of parallelism, advanced data management, etc,
- External (international) refereeing that focuses on the computational aspects of the projects, with the research priority handled separately, and with the excellence of the science and research team dealt with using the existing processes.

Long Term Implications—A National Merit Allocation Scheme

With the installation of the new Climate HPC facility in 2011-12, from funds provided under the Super Science initiative of the Commonwealth Government, there would be a three tier merit allocation scheme comprising:

- priority access for the Climate / Earth Systems Science research community,
- the Flagship Merit Allocation Scheme,
- the General Merit Allocation Scheme.

In fact, the top tier might be thought of as a dedicated, much larger flagship project which, by virtue of its size, would be allocated through a specialist, expert committee according to community agreed merit principles. The other national peak facility (at iVEC) under development through the Super Science initiative would be likely to have a similar allocation profile

In this environment, it is important to plan for a National Merit Allocation Scheme. It is natural to base this on the NCI Merit Allocation Scheme, given the considerable experience that has evolved in its implementation over many years.

Partner Access

Background

Access to NCI facilities, particularly the National Facility, other than through the Merit Allocation Scheme, is facilitated by the Partner and Affiliate Access schemes that are underpinned by the NCI Resource Allocation Model. The NCI Funding Agreement, and with it the original Business Plan, linked the concept of partnership (i.e., co-investment in the operations of NCI) with the governance framework.

Ultimately, this structure has not proven to be ideal for the long-term advancement of NCI and thus the two have been subsequently unlinked through the NCI Partner Service Agreement (which provides the services framework without any connection to the governance of NCI) and the proposal for a change in the Steering Committee composition (to strengthen the governance and strategic planning framework of NCI), put forward in the 2009-10 Business Plan and subsequently formalised through an exchange of letters and a Contract Variation in May 2009.

It is relevant that the original target of $1M p.a. for partner contributions has proven to be too ambitious for almost all research organisations, particularly in the light of the existing low base of co-investment in high-end computing within the sector. Accordingly, during 2008, the Steering Committee revised the definition of "partnership" to encompass just long-term commitment, adopting a modest threshold of $50K p.a. for partner membership from which to build co-investment in NCI. This, together with the leveraging of Commonwealth funds to make the partnership option more attractive, are the key concepts incorporated in the NCI Partner Service Agreement—a formal contract that extends over the three calendar years 2009-11, and which was signed by ANU and the initial partner, CSIRO, on 24 December 2008.

Partner Membership: organisations make long-term commitments to NCI over the term of the Partner Service Agreement (PSA) and, in return, receive access to a share of the facilities (determined using the NCI Resource Allocation Model described in Schedule 2 of the PSA), the value of which (in processor/core hours) increases with time as the system infrastructure is incremented.

Affiliate Membership: organisations acquire access to a fixed quota of resources (i.e., processor/core hours) for a fixed period under conditions and access pricing (prescribed in Schedule 2 of the PSA).

Partner Membership

Underpinning the definition of partner membership is NCI's mission to drive and extend Australia's uptake of high-end computing, and to provide a service that increases the scope and impact of Australian science.

To do so, NCI seeks to build mutually beneficial partnerships that increase the resources and services available to the Australian research community. From a base of Commonwealth (NCRIS) funding, NCI seeks to leverage additional investment from partner organisations which, through their commitment to NCI over the term of this agreement, provides them with a share of the computational resources that grows in absolute terms as the facilities are developed and incremented in capacity.

Since the cost structure of providing high-end computing facilities is not linear (due to the substantial fixed and quasi-fixed costs in both services and infrastructure), additional investments by partners not only increment the facilities but also reduce the cost of access (e.g., in $/core hour for compute time or $/GByte month for storage) to the benefit of all. In order to promote investment by partner organisations, partner contributions are weighted by a "partner multiplier factor" determined by the NCI Steering Committee (the present value of which is 1.1), thus leveraging Commonwealth funds in a mutually beneficial way.

Partnership membership requires a minimum annual investment, as determined by the NCI Steering Committee, that is to be sustained over the lifetime of this agreement, At this time, the minimum amount has been determined to be $50,000 annually, over the lifetime of the PSA.

While the threshold is modest, it provides an appropriate entry platform that allows research organisations and consortia to acquire resources with which to supplement grants under the MAS. From NCI's perspective, it provides a platform from which partners can build their reliance on, and subsequently increase their investment in, NCI. It is the latter that is seen as most significant, as is the certainty in cash flow that it provides NCI.

Affiliate Membership

Organisations which do not invest at least the threshold amount required for partner membership, or which do not invest for the sustained period over the lifetime of the PSA, are termed Affiliates of NCI and pay the Partner access rate for the period factored by a premium exceeding one that is determined, and revised from time to time, by the NCI Steering Committee. At the time of signing by the initial partners, the Affiliate access rate was set at 1.25 times the Partner access rate[7].

Present Status

The partner contributions from the two organisations anchoring NCI, ANU and CSIRO, have been secured through the Partner Service Agreement signed by their representatives on 24 December 2008. Respectively, they are contributing $3.4M ($1M cash and $2.4M in in-kind staff costs) and $3.3M ($3M cash and $0.3M in in-kind staff costs) annually over the three year lifetime of the PSA (2009-11).

Since then the Queensland (QCIF) and Western Australian (iVEC) consortia have signed the Partner Service Agreement, with each contributing at the $50K p.a. level.

Monash University has also signed the PSA, with its contribution in the first instance being $75K p.a.

[7] The Partner Access Rate (from which the Affiliate Rate is calculated) is determined as follows.
Capital costs (e.g., computers, power and cooling, and other infrastructure) are spread over the lifetime of the asset by dividing the cost by the asset lifetime. Recurrent costs (staff costs, facilities management, consumable power etc) are assigned to the period in which they are incurred. Combining the capital and recurrent costs for each period and dividing this by the total number of processor hours available in a six month period determines the cost of access for that period. Provision is made to accommodate variable contributions and contributions from late joining partners.

It is possible that at least some of the partners will increase their contributions in future years.

NCI has also been working with a number of leading research universities to build their uptake of NCI services. Some are undertaking internal reviews to determine the appropriate mix of high-end computing services for their future needs and we are hopeful of attracting some new university partners in coming years.

Intersect Consortium of NSW

Amongst NCI's successes over the past year has been in attracting the future partnership of the Intersect consortium of NSW.

Intersect Australia is the eResearch consortium of six NSW universities (Sydney, UNSW, Macquarie, UTS, Newcastle and Southern Cross) and SIRCA Ltd. During 2008, six NSW universities (UTS, Sydney, UNSW, Macquarie, Newcastle and Wollongong) applied to the ARC LIEF scheme for funding to update their local HPC infrastructure, and particularly to replace an ageing Dell Linux cluster. This proposal was subsequently successful with the six universities receiving some $500K from the ARC LIEF grant through the lead university, UTS. Since then, the six universities have confirmed their cash commitment of $390K to the ARC LIEF Grant, choosing not to pro-rate their contributions in line with the reduction in ARC funding from the requested $660K.

Earlier this year, the Intersect Board decided to consider two options:
(a) the direct renewal of local HPC infrastructure through the purchase of a new cluster system, and
(b) the acquisition of comprehensive services from NCI, based on the next generation National Facility system.

In February 2009, NCI was requested to prepare a comprehensive proposal for Intersect's consideration. This was sent to Intersect on 19 March 2009, immediately following the formal announcement of the signing of the contract for the Sun system, and was subsequently provided to Chief Investigators of the ARC LIEF grant and also to the Deputy Vice-Chancellors (Research) of the partner institutions for their consideration in late March 2009. At that time, Intersect commissioned an external consultant to undertake an independent assessment of the NCI proposal and also of what might be acquired in the way local facilities and services, within the same funding envelope.

The consultant's report, recommended that the consortium invest in NCI services, indicating that "the NCI proposal will most likely deliver the best outcomes for the delivery of high performance computing services to the NSW Universities and research teams". This report was provided to the CIs and DVCs (Research) in early May and it is understood that the CIs of the LIEF Grant were strongly in favour of adopting the NCI option. It is further understood that the Intersect Board, at its meeting of 12 May 2009, voted unanimously to accept this recommendation, put by the Intersect CEO, Dr Ian Gibson.

The NCI proposal adopted by Intersect comprises an up-front investment of $742K in infrastructure funds in 2009Q4, as part of the current Sun procurement, followed by equal payments of $375K per annum for a period of four years (2010-13). Annualising the infrastructure component, over the lifetime of the system (i.e., four years), leads to an effective annual investment of $560.5K p.a. It is this amount that determines Intersect's share of 4.17% according to the NCI Resource Allocation model.

Ongoing discussions with Intersect reveal their interest in committing at least part of the remaining $148K from the LIEF grant to increase the data storage on the National Facility that is available to NSW researchers. This will be pursued with Intersect during the development of the PSA Joining Agreement.

While all other NCI partners are committing funds for a fixed period (through until 31 December 2011), the Intersect consortium requires a four year agreement, corresponding to the lifetime of their investment of ARC infrastructure funds in the Sun Constellation system (2010-13). Accordingly, the draft Joining Agreement for Intersect provides for a special condition to accommodate this requirement, with the Steering Committee passing a recommendation (at its July 2009 Meeting), the effect of which is:

- for the Intersect Joining Agreement to the NCI Partner Service Agreement to apply until 31 December 2011 (the termination date of the PSA);
- the conditions applying to the Intersect/UTS partnership investment in NCI being rolled into any new partner service agreement, either with NCI represented by ANU, if NCI continues in its present form, or with any successor to NCI that continues to be represented by ANU from 1 January 2012.

  Should the NCI program be wound up, or any successor to NCI not be represented by ANU, then ongoing access to facilities and support services described in the agreement will continue to be provided by ANU as the host of the National Facility system, the service life of which, in the case of the Sun Constellation system, is January 2010 – December 2013.

Since both Intersect and UTS (on behalf of the ARC and partner universities) will be providing funds, the draft of the Joining Agreement prepared by NCI refers to a possible joint venture of Intersect and UTS. Both UTS and Intersect are presently considering their options as to how they wish to proceed, with the draft of the Joining Agreement being amended as appropriate to reflect their preferred structure.

At the time of preparing this Report, a letter to the ARC requesting its permission to repurpose the LIEF Grant Funds has been prepared and a the formal support of all partner institutions for this action is being sought.

3.3 Outreach

3.3.1 Raising the Profile of High-End Computing

NCI is most mindful of its mission of providing Australian researchers with access to world-class high-end computing services. Accordingly, the NCI Outreach program seeks to raise the profile and uptake of high-end computing in Australia and, in doing so, to make a difference to the impact and scale of Australian scientific achievements.

Over the past year, NCI has made good progress in engaging with a number of research communities, leading research institutions and NCRIS capability areas.

Of particular significance has been the work with the climate community through the Department of Climate Change, the Centre for Australian Weather and Climate Research (CAWCR) and the University Climate Consortium. This has included input into the development of the National Framework for Australian Climate Change Science, which ultimately informed the allocations made under the new Super Science Initiative, the implementation of support through the CT&T program for the climate science community, and the initiation of planning for the next generation system for climate change modelling involving a broad cross-section of the climate science, impact and adaptation community.

Also significant is NCI's work with the astronomy community through the HPC working group established by the NCRIS capability Astronomy Australia, and reported on under the CT&T Program.

Throughout the year, staff of NCI has engaged with a wide range of organisations to promote the role of NCI in the research community, to drive the uptake of high-end computing, and to further partnership opportunities where they might exist. These include the Bureau of Meteorology, CSIRO, Geoscience Australia, Monash University, the University of Melbourne, the University of NSW, University of Sydney, the University of Queensland, the Intersect Consortium of NSW, the iVEC consortium of Western Australia, the QCIF consortium of Queensland, the Australian Nuclear Science and Technology Organisation (ANSTO), the Australian Microanalysis and Microscopy Facility (AMMRF), and the Australian Research Council..

As noted previously, a number of these organisations (CSIRO, iVEC, QCIF, and Intersect) have confirmed partnership with NCI, and some others are considering opportunities to do so. The recent announcement of a substantial injection of funds for new HPC infrastructure through the Super Science initiative has changed the co-investment landscape considerably. From now on, the focus will change from one of collaboratively establishing infrastructure, from a low base, through to one of sustaining the new infrastructure and, particularly, in building the expertise base needed to maximise research outcomes.

NCI is also building its relationships with the ARC, with discussions held with the ARC CEO, Prof Margaret Sheil, and with a member of the ARC executive team, Dr Ian Mackinnon, attending a meeting of the NCI Steering Committee. Also of relevance is ARC's willingness to provide a representative on the NCI Merit Allocation Committee as an observer to assist in the development of our processes.

A forum of leading researchers intended to be held during the current reporting period has been postponed. It is intended that this be held either late in 2009 or early in 2010 to capture their experience of the new system, and to contribute to the planning processes for the next generation system (2012) by learning of the next wave of research drivers.

A User Survey will also be conducted in mid-2010 after the research community has had some time to experience the new peak system environment.

3.3.2 Promotional Material, Marketing and Website

Last year, following the establishment of NCI, a visual identity (emblem, logo etc) for use on letterhead, documents, websites etc was established and some promotional materials were developed. The www.nci.org.au website was put in place and the nf.nci.org.au was created by rebranding and updating the former nf.apac.edu.au site.

However, while the materials and web presence are functional, NCI is aware of the urgent need to upgrade its public face in order to highlight the range and quality of the services that are provided, and also the excellence of the research that is supported. The present web presence is somewhat rudimentary and the consistency between the two websites could be better. The gallery of research highlights is also quite dated, having been last updated in 2004 under APAC in the form of a broadsheet publication.

Spurred on by the launch of the new facility later this year, NCI has a number of initiatives underway that will revamp its public image, consistent with its mission of "providing a world-class high-end computing service to Australian researchers".

Website
A single new website is being developed that combines both the public face and technical user requirements that are implicit in the two existing websites. The opportunity is being taken to review the their content and to bring these under a single content management system.

Branding
In the process of specifying and developing the website, an opportunity has been taken to look at the overall branding of NCI and to update this as required. A concept has been sought from a design house and a fresh, dynamic image consistent with the mission of NCI and its "high-tech" services and infrastructure is evolving. Their brief is to develop the image (templates for webpages, printed materials, powerpoint presentation templates, business cards, and the like) while retaining the current logo, which links NCI to its antecedent, APAC, as a core element.

Research Highlights / Gallery
Cognisant of the excellence and impact of research supported by the National Facility, NCI understands the importance of highlighting this both generally, and more particularly to stakeholders, i.e., the sustaining organisations, potential investors and Government. During 2009, the NCI Office and staff of the National Facility identified a broad range projects for inclusion in the new research gallery (in both printed and web format). Some 15-20 projects were identified, with around eight of these being not only scientifically "high profile" but also "media friendly".

To this end, we are well underway with the process of updating the research highlights by engaging the prominent science journalist, Julian Cribb (www.sciencealert.com.au) to profile these projects and to develop written materials which can be used in both print and web formats. At the time of writing, eight vignettes of up to 750 words highlighting both the science and their implications/applications (as appropriate) have been prepared, with these being written in an accessible manner that might also attract the attention of the mainstream media. These vignettes have been developed from source materials provided by NCI (e.g., project proposals, reports, images etc) together with an interview with the project leaders. To further capitalise on the work, Mr Cribb has also developed short (150 word) "browsing summaries" of each project for inclusion on large posters which are being developed for exhibition at public events (conferences, exhibitions, facility launch etc).

Due to budget limitations, the remaining projects intended for our gallery, will be developed internally. Each year, however, we propose to increment the gallery through the inclusion of new projects as they arise through merit allocation applications or partner share allocations, with the intention of building up a strong, diverse and persuasive portfolio of research highlights.

3.3.4 Conference Activities

During the course of the reporting period, NCI participated in the eResearch Australasia 2008 conference. The National Facility staff hosted a booth and also offered a well attended half day workshop, while the Director participated in the eResearch Forum organised by the Department of Innovation.

For the eResearch Australasia 2009 conference, NCI will again be hosting a booth, while the Director is serving as a member of the conference Program Committee.

3.3.5 Training Programs

The NCI National Facility continues its outreach program in which staff of the NF provide user workshops and training visits to research communities and research organisations as and when required. Details of these can be found in the National Facility Reports (Appendices 3 and 4).

4. GOVERNANCE AND MANAGEMENT

Background and Current Status

The Funding Agreement specifies that NCI will be governed by ANU on advice from a Steering Committee (SC) with respect to the following matters:

- strategic plans for the national computational infrastructure and services,
- business and marketing plans for the Program,
- project plans for outreach activity,
- budget allocations within the Program,
- the general direction of implementation of the Program and associated delivery of services to users, and
- arrangements for promotion, collaboration and cooperation among the contributors, counterparts, and research organisations nationally and internationally.

In its 2007-08 Annual Report, NCI foreshadowed strengthened strategic planning processes for HPC services through the establishment of a national high-end computing council. This underlying concept was subsequently formalised in the 2009-10 Business Plan with a proposal to build this upon the foundation of the NCI Steering Committee, a body which draws together organisations with a strong commitment to, and a financial stake in, the advancement of the nation's high-end computational infrastructure.

The 2009-10 Business Plan also proposed the formalisation of the composition of the Steering Committee[8]. Following an exchange of letters between NCI and the Department of Innovation in early May, a revised composition of the Steering Committee which confirmed the existing membership (including the previously designated "potential partners") and provided for representation of the University of Melbourne (representing the interests of the Victorian Life Sciences Computation Initiative) was confirmed in a Deed of Variation signed on 26 May 2009.

The current composition of the Steering Committee, as agreed in the Deed of Variation, is:

- An independent Chair (presently Emeritus Professor Mark Wainwright, AM)
- Director, NCI (presently Professor Lindsay Botten)
- Representatives of:
  - The Australian National University (presently Professor Robin Stanton, Pro Vice-Chancellor)
  - CSIRO (presently Dr Alex Zelinsky, Director, CSIRO ICT Centre)
  - Bureau of Meteorology (presently Dr Neville Smith, Acting Director)
  - Geoscience Australia (presently Dr Chris Pigram, Deputy CEO)
  - University of Melbourne (Professor John O'Callaghan, Interim Director, VLSCI, from 2009Q3)
  - Research Intensive Universities (presently Prof Doug McEachern, Deputy Vice-Chancellor (Research and Innovation), University of Western Australia)
- Chair, NCI Merit Allocation Committee (presently Professor Brian Yates, University of Tasmania)

Ongoing Role of the NCI Steering Committee

Following discussions between the Steering Committee Chair and the Director from NCI, and Ms Anne-Marie Lansdown and Mr David Luchetti from the Department of Innovation, that were held on 3 July 2009, together with subsequent correspondence (dated 10 July and 16 July 2009), the role of NCI in providing strategic advice to Government on national high-end computing services was clarified. The Department, in its letter of 16 July 2009 from Mr Lansdown to the Chair of the Steering Committee, noted that such matters would be referred to the newly constituted NRIC and that the Department was considering "proposing a variation to the funding agreement to remove the obligation on ANU to provide strategic advice as envisaged in the NCI funding agreement". NCI in its reply of 24 July from the Steering Committee Chair noted that "when NRIC chooses to seek such advice, NCI will be ready and willing to contribute in whatever way [it] can".

[8] At the time of the submission of the 2009-10 Business Plan, the Steering Committee comprised representation from ANU (as the contracting organisation), CSIRO (as a signed partner of NCI), representatives of a number of potential partners, as well as the ex-officio positions of an independent Chair, the NCI Director and the Merit Allocation Committee Chair. Following the submission of the 2007-08 Annual Plan and subsequent correspondence between NCI and the Department of Innovation, the terms of the potential partners had been extended until 30 June 2009.

5. PERFORMANCE INDICATORS AND MILESTONES

The NCI Business Plan 2007-11 listed milestones for the July 2007 to June 2009 period. These, together with an outline of progress against them, are detailed below.

Progress against Milestones—January 2008 to June 2009

Activity	Expected Date	Progress against milestone
Appointment of Director	Q2 2008	The substantive Director of NCI commenced duties on 26 May 2008.
Start procurement process for new peak system	Q2 2008	A joint tender between the Bureau of Meteorology and the ANU was let on 3 April 2008 and closed on 29 May 2008.
Call for EoI for Specialised Facilities	Q2 2008	The call for the Specialised Facilities Expressions of Interest was issued on 21 July 2008 and closed on 20 August 2008. Agreements for two Specialised Facilities (in Bioinformatics (UQ) and Imaging and Visualisation (Monash) are close to finalisation for commencement in January 2010. Each has funding of $1.2M (over two years) for services being acquired for three years 2010-12. Progress for this milestone has been acknowledged by the Department of Innovation in the early payment (June-July 2009) of the FY2009-10 tranche of funds.
Call for EoI for Computational Tools and Techniques	Q2 2008	The Steering Committee decided against a separate call for expressions of interest, instead recommending that NCI proceed with support for two communities, Earth Systems Science and Astronomy, identified in the Expressions of Interest for the SF Program. Support for data intensive computation through cloud computing was proposed in the 2009-10 Business Plan that was accepted by the Department of Innovation in May 2009. Progress for this milestone has been acknowledged by the Department of Innovation in the early payment (July-August 2009) of the FY2009-10 tranche of funds.
Transitional system installed	Q2/Q3 2008	The tender was let in June 2008, with an order placed in September 2008. The system was installed in December 2008 and commissioned in January 2009.
Participate in eResearch Conference; User Forum	September 2008	Participation in the eResearch Conference and AeRIC Forum took place in September/October 2008.
MAS Call	October 2008	The MAS processes were revised in Q3 2008. The MAS call under these new process, using an extensively revised web-based application form, occurred in October 2008.
Attend SC08 Conference	November 2008	SC08 was not attended due to its clash with the commencement of contract negotiations for the new Sun system.
MAC Meeting	December 2008	The MAC meeting took place on 26 November 2008.
Contract for the new system signed	Q4 2008	The contract with Sun Microsystems for the new system was signed on 11 March 2008.
Deploy Stage 1 of the new system	Q2 2009	The general user service from the first phase was available from 17 September 2009.

The performance of NCI against agreed performance indicators is provided in Attachment 1.

6. FINANCES AND RESOURCES

6.1 Statement of Income and Expenditure

The ANU Certified Income and Expenditure Statements for the period 1 July 2008 to 30 June 2009, is provided at Attachment 2.

The expenditure reported in Attachment 2 relates to operations of the NCI Office over the entire period (associated with the Planning, Policy and Access functions of NCI), capital outlays for the new developmental system at the National Facility, and operational costs of the National Facility for the period (January-June 2009). As is outlined in the 2008-09 and 2009-10 Business Plans, provision for the operational costs of the National Facility through until the end of 2008 was made by the Board of APAC[9] and is shown in the table immediately below.

The Facilities Management Agreement was renegotiated with ANU in late 2008, and approved by the Steering Committee at its November 2008 meeting in order to update the previous agreement negotiated under APAC in 2006.

National Facility Operational Expenses	July-December 2008	January-June 2009
	From former APAC funds	From NCRIS funds
Facilities Management	$ 257,500.00	$ 367,500.00
Data Services	$ 266,500.00	$ 287,500.00
Consumable Power	$ 238,971.15	$ 121,984.78
Software	$ 61,759.00	$ 31,553.40
Staff (In-kind)	$ 1,200,000.00	$ 1,200,000.00
Total	$ 2,024,730.15	$ 2,008,538.18

The consumable power that is listed for the period July-December 2008 is the actual expense that was incurred. The corresponding listing for the period January-June 2009 shows actual expenses for the period January to March 2009. Electricity expenses for the remaining three months (April-June 2008) will rollover into next year's report due to delays in billing and reporting.

Staff services within the National Facility are borne by the annual in-kind contributions of $2.4M from ANU, as detailed in the Budget presented in the Business Plan.

No expenditure was incurred for the Specialised Facilities and Computational Tools and Techniques programs as these had not been initiated during the reporting period.

6.2 Cash and In-kind Contributions

The Financial Statement in Attachment 2 details the partner income which to be interpreted as follows:

- ANU contributed $500K for the period January-June 2009, consistent with the Partner Service Agreement and the 2009-10 Budget.
- CSIRO contributed $1,500K for the period January-June 2009, consistent with the Partner Service Agreement and the 2009-10 Budget. The additional amount of $40K is for its share during the transition phase in 2008.
- The amounts of $40K listed for each of iVEC and QCIF are for shares in July-December 2008 and an ongoing amount in 2009 prior to their conversion to partnership at the $50K p.a. rate. The difference will be reflected in next year's report due to delays in billing and receipting.
- The additional amount of $150K attributed to CAWCR is a dedicated investment by CSIRO in the development system to provide urgently needed additional access for climate model development during 2009.

During the reporting period, the in-kind contribution of $2.4M was provided by ANU through the provision of 15.5 EFT in staff support for the National Facility. Through the partner service agreement, CSIRO also contributes $300K p.a., with the value of this during the reporting period being $150K (January – June 2009 following the signing of the partner share agreement).

[9] This is highlighted in the blue typeface entries of the Budget on page 24 of the 2009-10 NCI Business Plan.

Attachment 1: Report on Performance against Key Indicators

The NCI Agreement and Business Plan list key objectives, performance indicators and outcomes. The report on the performance of NCI against these indicators for the period July 2008 to June 2009 is given below.

Objective: Develop a national strategy for advanced computing to support eResearch in Australia.

A national advanced computing strategy to support eResearch is in place and maintained.

- NCI builds on the successful National Facility program of its antecedent APAC, informed by its mission of "providing Australian researchers with world-class high-end computing services".
- NCI is building strong partnerships sustained by a substantial base of co-funding with leading research organisations that establish it as a service integral to the advancement of cutting edge science in Australia.
- Partnerships informed by particular science drivers (particularly earth systems science / climate modelling) are being established (with BoM, CSIRO and research intensive universities).
- A framework for advancing national high-end computing strategy through growing engagement with leading national research organisations and universities, underpinned by co-investment and co-operation was foreshadowed in the previous annual report and proposed in the 2009-10 Business Plan. However, from the Department's perspective (as communicated in the 16 July 2009 letter from Ms Lansdown to the NCI Steering Committee Chair) the recent announcement of new investments in this space means that the need for future strategic advice is not pressing and that such matters would be dealt with in future through the recently constituted National Research Infrastructure Council (NRIC).
- NCI continues to work with its key communities, particularly the climate research community, planning for their long-term strategic needs.
- NCI will also evolve the Merit Allocation Scheme in the national interest and prepare for the introduction of a National Merit Allocation which it is uniquely well placed to host.

Demonstrated progress towards the implementation of the strategy.

- The NCI National Facility has consistently provided a world-class high-end computing service for Australian researchers, and the new peak system (Sun Constellation) will provide a major boost to high-end users in this country. Indeed, for the first time in two or three years, Australia will have a machine that is internationally significant. At the time of writing, with the first phase of the new facility having been commissioned for only a few days, already there is strong acclamation from the user community about the very favourable performance increases and the ease and simplicity of the transition process from the former system.
- The development/transition system has also provided a significant performance boost during 2009 and has proven to be an excellent development platform for the research community in preparation for the new Sun Constellation system.
- Resource allocation strategies that encourage partner investment continue to evolve.
- Merit allocation processes continue to be strengthened with significant efforts now being directed towards priority research under the flagship program.

Objective: Strengthen relationships to provide national cooperation on Australia's advanced computing infrastructure.

Membership and participation in NCI categorised by degree of involvement and contribution of organisations.

- Present commitments for co-investment in NCI (directed towards the National Facility) are approximately $6.85 M p.a. (with major contributions of $3.4M p.a. from ANU and $3.3M from CSIRO) rising to slightly in excess of $7.4M in 2010 following Intersect joining the NCI partnership. This is to be compared with the annualised contribution of $6.5M from the Commonwealth under NCRIS. There will also be strong levels of coinvestment in the Specialised Facilities, but these are yet to be confirmed.
- Work is ongoing to attract additional co-investment in the National Facility from leading research universities and other leading research organisations.
- An extensive program of visits to potential partners has continued during the reporting period. The primary aim of these visits is to build and extend relationships (co-investment, collaboration and co-operation) with leading researchers, high impact research groups, and senior managers of these organisations in a way that extends the range of services offered by NCI and which sustains its operations. A list of organisations visited is presented elsewhere in the body of this report.
- NCI continues to work closely with the climate community (BoM, CSIRO, CAWCR, UCC) in assisting both the planning for new facilities and the uptake of services by this community. Support under the CT&T program will be particularly significant in 2010 in the lead up to the installation of the new climate HPC facility (funded through the Super Science program).
- A forum of leading researchers, to be held following the commissioning of the new peak system, will be held either late in 2009 or early in 2010.

Funding received from sources other than NCRIS.

- Present commitments for co-investment in NCI exceed $6.8M p.a.

Objective: Improve the peak computing capabilities of the NCI National Facility to serve the demands of Australian researchers.

Key research achievements obtained by users of the National Facility.

- A summary of all NCI (MAS) projects including brief descriptions is also available at http://nf.nci.org.au/accounts/grants/.
- The Lead Chief Investigators of these projects are required to submit an annual report on their projects. The reports for 2008 and previous years are available on-line at http://nf.nci.org.au/annual_reports/.
- National Facility staff have developed publicity materials (posters and brochures) highlighting the excellence of the research conducted on NF systems. These materials are now being substantially upgraded with assistance from a leading science writer and will be used in the new Research Highlights / Gallery on the NCI website, and also in printed form.

Significant developments in computational tools and techniques on the NCI and partner facilities.

- The Computational Tools and Techniques program has a different orientation to that formerly conducted within APAC. Under NCI's program, targeted support will be provided to a small number of research areas (climate science, astronomy) and new services will be developed for data intensive computation using cloud computing.
- Cash flow constraints (a combination of the cash flow to NCI from the Commonwealth plus the challenging economic environment in which the new peak facility has been acquired) have necessitated delaying the start of the bulk of this program until 2010.
- The project in data intensive computation is underway and project plans for the climate science and astronomy communities are being finalised for implementation in 2010. Staffing agreements have also been developed.

Capability of the National Facility relative to countries of similar size and development such as Canada, Sweden and Korea.

- Based on the June 2009 release of the TOP500 list (www.top500.org) for June 2009, Australia had only one system (from the digital media industry) within the top 500 list, ranked at position 279. The current National Facility has fallen off the TOP500 list, last appearing at rank 200 in the November 2007 release of the list.
- The decline in Australia's representation on the TOP500 list is demonstrated in the table below, commencing with June 2005, the first period in which the current National Facility was recorded in the list.

List Date	Australian Systems in TOP500 List	Australian Aggregate Performance (ΣRmax Gflops)	Highest Ranked Australian System	International Rank of National Facility	Australian Rank of National Facility
June 2005	5	19,592	26	26	1
November 2005	11	34,460	36	36	1
June 2006	9	31,023	47	47	1
November 2006	4	20,670	71	71	1
June 2007	4	22,522	107	107	1
November 2007	1	8,974	200	200	1
June 2008	1	9,239	468	-	1
November 2008	1	21,906	145	-	1
June 2009	1	21,906	279	-	1

- Had the new Sun Constellation system been fully commissioned in time for the June 2009 list, it would have ranked (approximately) in position 24.
- In June 2008, two of our traditional comparator nations of Canada, Korea and Sweden had research systems listed in the TOP500.

  Canada had 8 systems in the list at ranks 16(U), 80(U), 163(U), 224(U), 254(I), 292(U), 295(G), 317(I), where the U/I/G qualifiers indicate their location in university/research, industry or government environments. This is a significant increase over the previous year in which Canada reported only two systems at ranks 249(G) and 395(U).

  There are no systems listed for Korea in the June 2009 list.

  Sweden lists 10 systems in the TOP500 at ranks 30(G), 83(U), 88(U), 176(I), 221(I), 293(I), 350(I), 460(I), 477(I), 494(I).

Capability of the National Facility relative to the needs of the Australian research community as indicated by extent and type of demand.

- The demand for the National Facility is demonstrated by:
  - The large number of projects, principal investigators and associated researchers supported by the National Facility.
  - The geographic distribution of users for the MAS (coming from 24 research organisations in every State and the ACT).

- The broad base of application areas and research disciplines.
  These are detailed in the Tables of Section 3.1.1 and in the National Facility Operations reports in Appendices 3 and 4.

- The demand for MAS grants in the December 2008 round exceeded the available resources by a factor of approximately 2.2. In recent years, this figure has saturated as researchers become aware that there is no point in applying for resources which are not available. However, this oversubscription ratio spiked this year, presumably in anticipation of the new peak system.
- The resources available under MAS allocated at the December 2008 round constituted 46% of the available facilities. This contrasts with the 73% of the facility allocated under MAS in the December 2007 round but was in line with the 42% share allocated during 2004-06 under APAC. The share available to MAS allocations for 2008 was an anomaly corresponding to a minimum in partner investment in the transition between APAC and NCI. The MAS allocations for 2009 reflect the shares allocated to partners (including the two major partners, ANU and CSIRO).
- The increased demand for data storage and data management services is following a similar trend to the computing demands.
- The quality of the research projects supported is indicated by the high proportion of grants which are associated with ARC and external support.
- Further details on the demand for the National Facility are given in Section 3.1.1 and in the operations reports in Appendices 3 and 4.

Extent of use of the National Facility indicated by the number of States represented by users, projects, users and resource allocations.

- During the six month period from January-June 2009, there were 228 projects through the MAS and partner shares associated with 977 researchers with login accounts (and a further 65 researchers without accounts).
- The MAS scheme accounted for 153 of the 228 projects, of which 178 were supported by research funding (ARC, industry, or other).
- Further details on the extent of use of the National Facility are given in Section 3.1.1 and in the National Facility reports (Appendices 3 and 4)

Availability and performance of the National Facility indicated by the available system units, used system units, efficiency of operation and extent of parallel computation on the National Facility.

- During the six-months from January to June 2009, the SGI Altix 3700 system was available for 98.16% of the total time theoretically available. For the new SGI Altix XE, the figures are 72.18% availability in 2009Q1 (when the system was still being commissioned and minor problems dealt with) and 99.13% in 2009Q2 after the facility had been bedded down.
- Overall system utilisation for the period has been 96% which is extremely high and in the ranks of international best practice, nearing the capacity of the system for the current job mix. During this time the MAS share was been fully used. The comparison purposes, the corresponding system utilisation for the July-December 2008 period was 97.4%.
- Workloads have continued to be very high, although the solution to this is in sight with the commissioning of the new Sun Constellation.
- The percentage of projects using the parallel capability of the peak system is around 92% of the available resources of the system. This has increased from 85% in 2007-08 and 75% reported in 2006-07. This percentage accounts for parallel jobs utilising the high speed interconnect of the peak system. Approximately two-thirds of the usage uses 16 cpus or more, while 6% of the total utilisation uses about one-eighth of the system of more—a figure that is somewhat lower than that seen on international peak systems. This matter is considered further in the report under the heading of a Flagship MAS.
- All jobs utilise the global filesystem, application packages and other software. The performance scales uniformly across the entire system and resource allocation is based on underlying hardware resources rather than dedicated nodes for specific (software) purposes.
- Further details on the availability and performance of the National Facility are given in Section 3.1.1 and in the National Facility Reports (Appendices 3 and 4).

Summary of surveys of users of the National Facility showing their level of satisfaction.

- The Director and the Manager of the National Facility work closely in developing plans for the ongoing development of NCI systems and services, and in reviewing operational issues that relate to the experience of users.
- A customer survey was last conducted in March 2006 to gain insight into the level of user satisfaction following the installation of the new SGI system (installed in 2005). Then, the overall customer satisfaction was at a very high level (4.6 out of a possible 5) indicating that the National Facility is well supported by its users.
- The customer / user survey planned for 2009 has been postponed until 2010, following the commissioning of the new peak system. Although it had been intended to conduct a survey in 2009, the introduction of the new development / transition system would not have materially changed the users' experience of NCI and so it was decided to delay the survey until the commissioning of the new peak facility.
- The transition to the new facility has led to a number of very favourable comments from users (about the quality of service and the performance of the new facilities) being received. Accordingly, NCI looks forward to the 2010 user survey once the new system is completely installed and bedded down.

Objective: Implement specialised computing capabilities through NCI Specialised Facilities to serve the demands of Australian researchers.

Key research achievements obtained by users of Specialised Facilities.

Capability of Specialised Facilities relative to the needs of the Australian research community as indicated by extend and type of demand.

Extent of use of Specialised Facilities indicated by the number of States represented by users, projects, users and resource allocations.

Availability and performance of the Specialised Facilities is comparable to the National Facility.

Summary of surveys of users of Specialised Facilities showing their level of satisfaction.

- The Specialised Facilities Program will commence in early 2010 with allocations made under the 2009 round of Merit Allocation. Accordingly there is no progress to report against these indicators.
- Cash flow constraints (a combination of the cash flow to NCI from the Commonwealth plus the challenging economic environment in which the new peak facility has been acquired) have necessitated delaying the start of this program until 2010.

Objectives: (1) Develop and deliver outreach services for users of advanced computing infrastructure. (2) Participation and contribution of organisations in the Outreach program. (3) Number and types of outreach activities conducted.

Extent and nature of working linkages between industry and research organisations.

- NCI continues to work with research communities, with a particular focus on the climate research community, to build the uptake of high-end computing services.
- The substantial investment of ANU, CSIRO, and Intersect (from 2010) shows strong commitment to increasing their use of high-end computing.
- NCI will be co-funding two positions in each of Melbourne and Sydney (to be hosted through Monash University and Intersect respectively). These positions, which will commence from the beginning of 2010, will be embedded in the National Facility Support Team and will assist in NCI's outreach activities in promoting the uptake of high-end computing services.
- The Specialised Facilities and Computational Tools and Techniques Program extend NCI's outreach to the bioinformatics, characterisation, and climate research and astronomy communities.

Increased skills of users of advanced computing infrastructure.

- National Facility workshops have continued to be run.
- Additional support for targeted user areas is planned through the Computational Tools and Techniques Program commencing in 2010.

Appendix 1: Members of the APAC Merit Allocation Committee

The following is a list of the current members of the Merit Allocation Committee as at 30 June 2009.

Professor Brian Yates (Chair)
School of Chemistry
University of Tasmania

Professor Julian Gale
Nanochemistry Research Institute
Curtin University of Technology

Prof Geoff Bicknell
Research School of Astronomy and Astrophysics
The Australian National University

Associate Professor Jim Denier
Department of Applied Mathematics
The University of Adelaide

Associate Professor Louis Moresi
Department of Computational Mathematics and Geophysics
Monash University

Professor Alan Mark
Institute for Molecular Bioscience
University of Queensland

Professor Ross McPhedran
School of Physics
University of Sydney

Dr Tony Hirst
Marine and Atmospheric Research
CSIRO

Dr Alf Ulhlerr
Advanced Scientific Computing, CSIRO ICT Centre
CSIRO

During the period, Dr Ben Evans, NCI National Manager served as an advisor to the committee while Ms Judy Jenkinson from the National Facility provided the committee with secretariat services.

Appendix 2: Members of the NCI Steering Committee

The members of the NCI Steering Committee at 30 June 2009 were:

Emeritus Professor Mark S. Wainwright AM (Chair) The University of New South Wales SYDNEY NSW	Independent Member
Professor Robin Stanton Pro Vice-Chancellor The Australian National University CANBERRA ACT	ANU
Professor Brian Yates University of Tasmania HOBART TAS	Independent Member and Chair, Merit Allocation Committee
Professor Lindsay Botten Director, NCI The Australian National University CANBERRA ACT	NCI

Partners and prospective partners represented on the NCI Steering Committee at 30 June 2009 were:

Dr Alex Zelinsky Director, CSIRO ICT Centre SYDNEY NSW	CSIRO
Dr Neville Smith A/Deputy Director (Research & Systems) Bureau of Meteorology DOCKLANDS VIC	Bureau of Meteorology
Dr Chris Pigram Deputy CEO and Chief of Geospatial and Earth Monitoring Division Geoscience Australia CANBERRA ACT	Geoscience Australia
Professor Doug McEachern Deputy Vice-Chancellor University of Western Australia PERTH WA	Deputy Vice-Chancellors (Research)

Observers invited to attend NCI Steering Committee meetings at 30 June 2009 were:

Professor Leon Stirling MELBOURNE VIC	University of Melbourne
Dr Ian MacKinnon (March 2009 meeting only) CANBERRA ACT	ARC (Executive Director)
Dr Ben Evans CANBERRA ACT	ANU/ NCI National Facility Manager
Mr David Toll Canberra ACT	CSIRO (Standing Alternate for Dr Zelinsky) Director, Property and Information Services

Appendix 3: National Facility Operations Report – July 2008 to December 2008

This Appendix outlines the systems in the National Facility and the systems and user support provided by ANU staff working in the National Facility during the period July 2008 to December 2008.

Merit and time Allocations

Merit Allocation Scheme (MAS) for 2009

The Merit Allocation Committee met in Canberra on November 26, 2008 to finalise the assessments and to grant resources to projects for 2009. The total MAS share for 2009 was 46% of the total resources, representing approximately 12.6 MSU for the year. The MAS share provides time on both the existing peak SGI AC computer on the new development SGI Xe cluster, and associated storage for computational output, and hosting critical datasets required for computational workflows.

MAS Committee Time Allocation Summary

The Merit Committee Chairman has provided a detailed report to the Steering Committee. In summary, the request to the MAS was 28.4 MSU (2.2 times the share), and the total grants awarded totalled 15.3MSU (1.2 times the MAS share). There were 163 applications for time including 33 new projects. Of these, 151 projects were awarded time of which 24 were new MAS projects starting in 2009. There were 2 new start-up projects. The time granted for 2009 as a percentage of the overall MAS share was:

	2006	2007	2008	2009
Top 10 projects	42%	40%	45%	42%
Top 20 projects	60%	58%	62%	60%
Top 40 projects	78%	80%	82%	80%

The request and allocation to the MAS based on major ARC category is shown below.

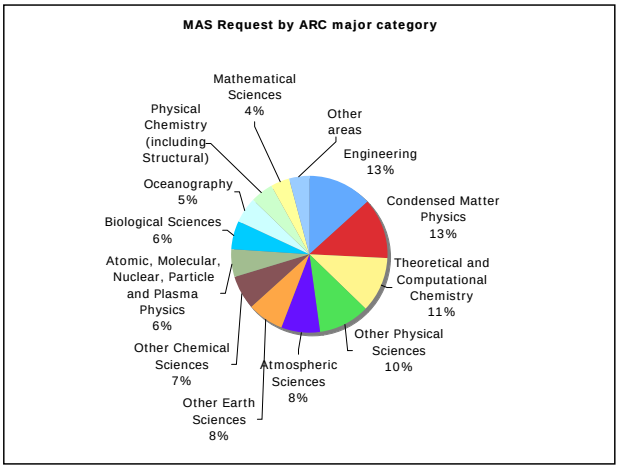

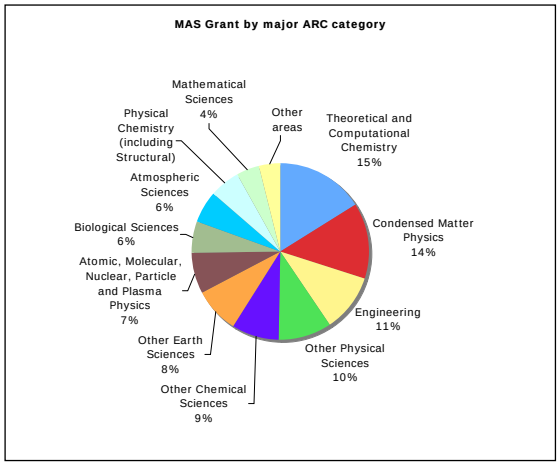

Overall Request, Grant and Success by Institution

The following table represents the overall outcomes across the research organisations using the facility for combined access to MAS and others shares:

University	Request (SU)	Grant (SU)	Success
ANU	2187683	1764430	81%
CSIRO	1026619	818744	80%
Sydney	965290	736349	76%
Monash	906890	428337	47%
UQ	616139	319844	52%
Melbourne	568500	302748	53%
UNSW	626050	291224	47%
Adelaide	425300	284372	67%
Curtin	295000	257750	87%

RMIT	335500	225412	67%
UTS	139330	73500	53%
UTAS	234750	64582	28%
UWA	103750	63624	61%
Newcastle	195000	50000	26%
Murdoch	72500	37500	52%
La Trobe	45000	30000	67%
ADFA	52402	23258	44%
Other Australian Research Institute	144500	22000	15%
Griffith	18200	12430	68%
Other Queensland	54666	9333	17%
James Cook	15000	8571	57%
Macquarie	12165	8000	66%
Wollongong	6000	5000	83%
Swinburne	10000	4274	43%
Southern Queensland	10000	4000	40%
Australian Gov Dept	3500	3500	100%
South Australia	3000	3000	100%
Other NSW	2500	1000	40%

2009 Projects and Users

The MAS and shareholders granted a total of 217 computational projects through MAS and other shares with 919 researchers. There were an additional 64 researchers associated with projects but who did not have a login-names, bringing the total number of researchers involved with the projects to 983.

Usage Data for 2008H2

In the previous period, the National Facility serviced a total of 223 computational projects through MAS and other shares with 992 researchers. There were an additional 73 researchers associated with projects but who did not have a login-names, bringing the total number of researchers involved with the projects to 1065. There was only one new MAS project awarded time starting in 2008H2 and four new start-up projects.

The tables below present the overall number of projects and researchers and their grants and usage during the period broken down by projects under the MAS and other shares. More details on the individual projects are available on the NF website. It should be noted that some researchers have access to both MAS and another share project. Projects may also have more than one source of grant.

Of these projects, 157 had support from ARC or external grants. (Note that CSIRO projects are not eligible for ARC grants).

The MAS was 1.58 times over-requested at the start of the period 2008H2. The committee granted 1.2 times the formal MAS share.

Reports on Individual Projects

Principal Investigators of Merit Allocation Scheme projects are required to submit a report on the project each year. These reports are available on-line at http://nf.nci.org.au/annual_reports/.

A list of all MAS projects including brief descriptions can also be found at http://nf.nci.org.au/accounts/grants/

Overall Usage of Computational Time

Division	No. of PIs	No. of Projects	No. of Researchers with login-names	No. of Researchers (total)	No. with ARC Support	No. with NHMRC Support	No. with NCRIS Support	No. with Industry Grant	No. with Other Funding	Grant as % of Total Resource Available	Usage as % of Total Resource Used
MAS	148	148	775	831	86	4	1	9	34	84.2	74.2
Other Shares	72	77	357	378	30	3	2	3	12	30.9	24.9
Start-up	9	9	15	18	0	0	0	0	0	<0.1	<0.1
TOTALS	214	234	1016	1092	116	7	3	12	45	115.2	100

Other Shares

Division	No. of PIs	No. of Projects	No. of Users with login-names	No. of Researchers (total)	No. with ARC Support	No. with NHMRC Support	No. with NCRIS Support	No. with Industry Grant	No. with Other Funding	Grant as % of Total Resource Available	Usage as % of Total Resource Used	Formal Percentage Share of the System
ANU	45	50	267	280	26	3	1	3	8	28.5	23.6	25.00
CSIRO	14	14	47	51	1	0	1	0	5	1.2	1.1	1.15
IVEC	4	4	12	13	0	0	0	0	0	0.5	0.1	0.46
QCIF	5	5	12	14	2	0	0	0	0	0.4	0.1	0.46
Other Shares	2	2	9	10	0	0	0	0	0	0.3	<0.1	<0.01

Note: Totals represent unique PI, researchers, and users.

Overall Usage and demand for the System

While there is a shortage of cycles and hence reasonably long queues, the ANU PBS scheduler is adaptive and intelligent and thus users overall still get reasonable access to the system. The following graphs show the amount of queued work on the system in the second half of 2008. There has been approximately 3.6 times the number of CPUs queued on the system than available, rising to a peak of 5 times oversubscribed.

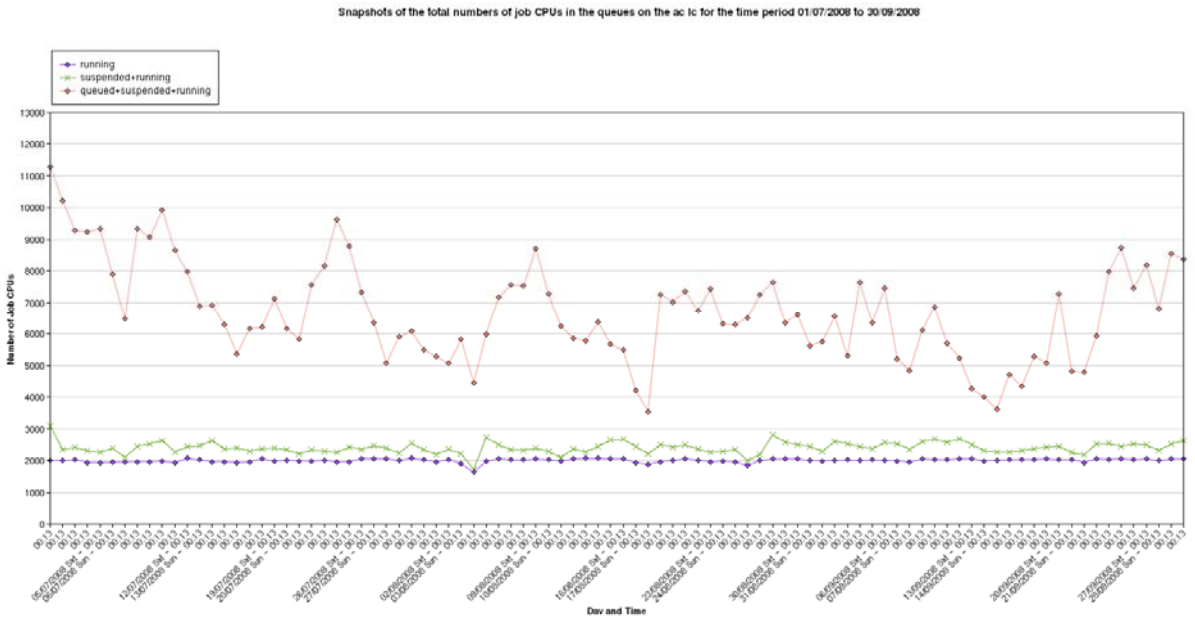

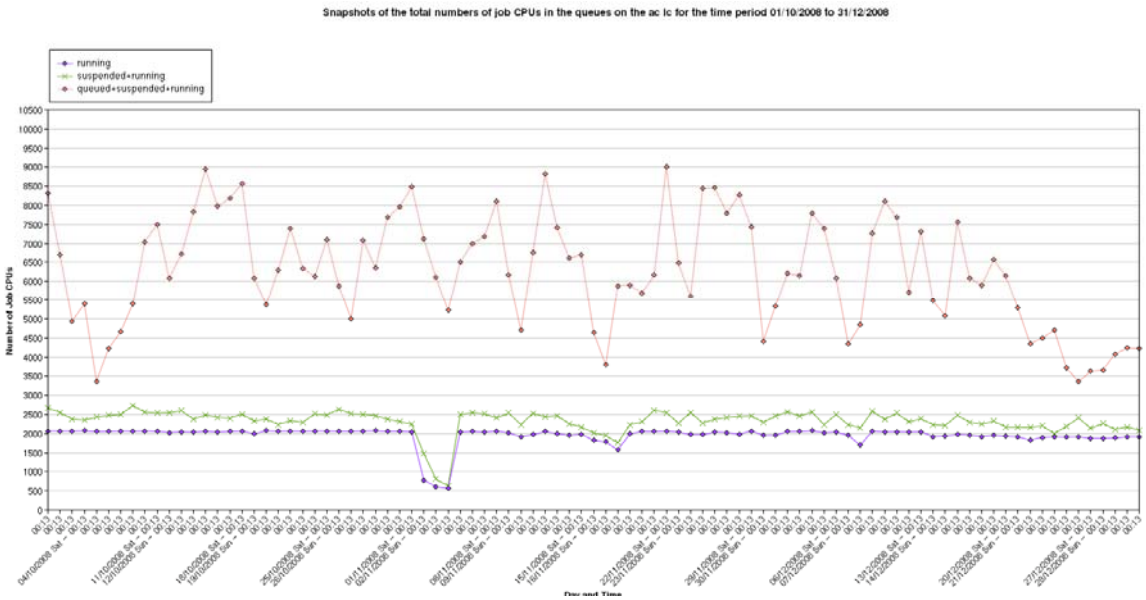

The MAS share has been fully used on the system over the whole period as shown in the following graphs. Overall, system utilisation was an average of 97.4% utilisation for the period, which is extremely high, and is nearing to capacity of the system for the current job mix. The remaining cycles could potentially be used with more small single CPU jobs.

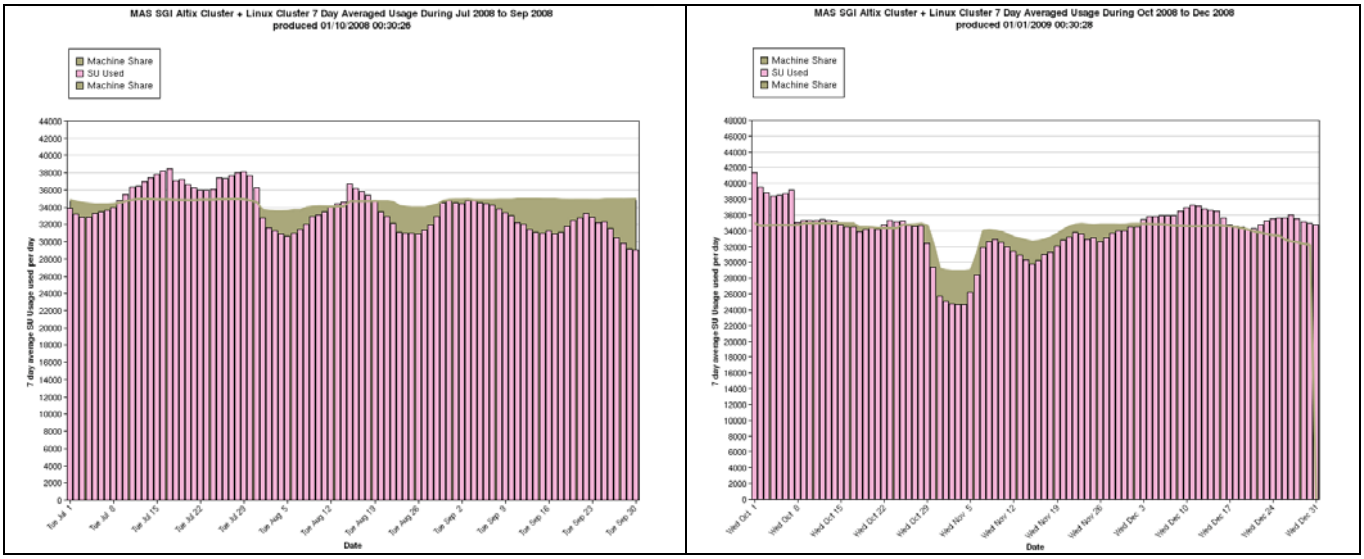

The following graphs show the demand for MAS share over a longer period; 2004 to present. The graph clearly shows the increased demand for peak computing cycles for 2009.

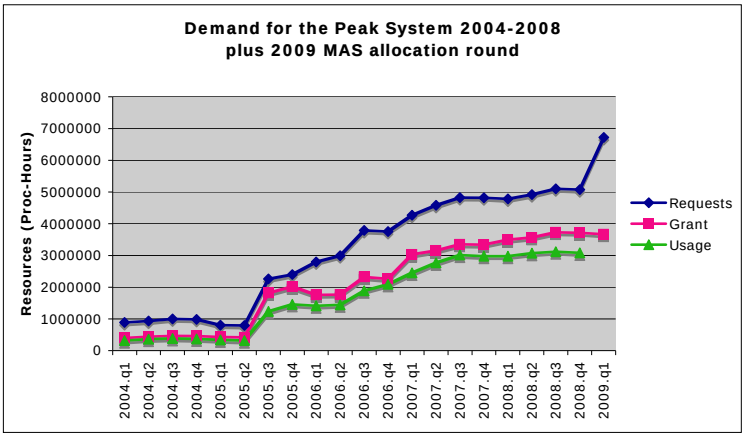

Usage by Research Field

A summary of the usage by research disciplines during the period is given in the two pie-charts below. The first is restricted to MAS computational projects only while the second shows usage over the entire user base.

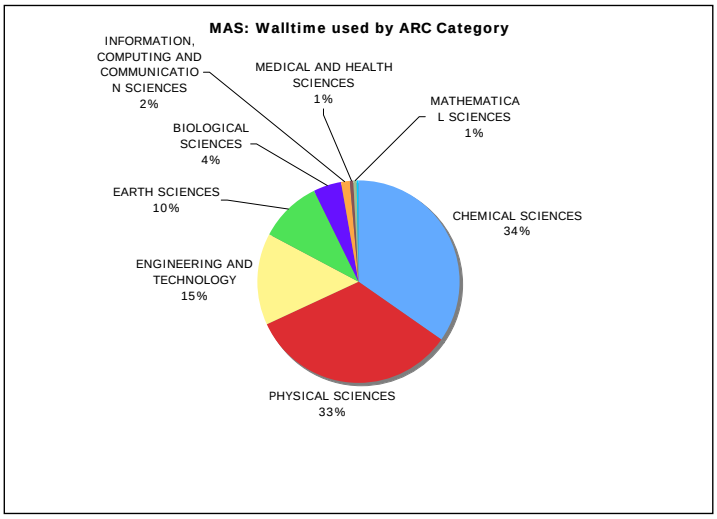

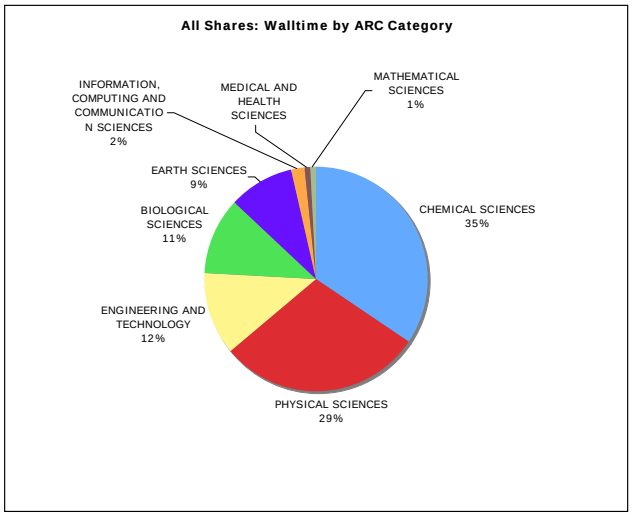

Measure of Parallel Use of the AC

The following graph indicates the percentage of the ac system resources used by parallel jobs. The NF encourages users who might otherwise run an ensemble of tasks as a parallel task, to run the jobs serially to allow the scheduler to make maximum use of the system.

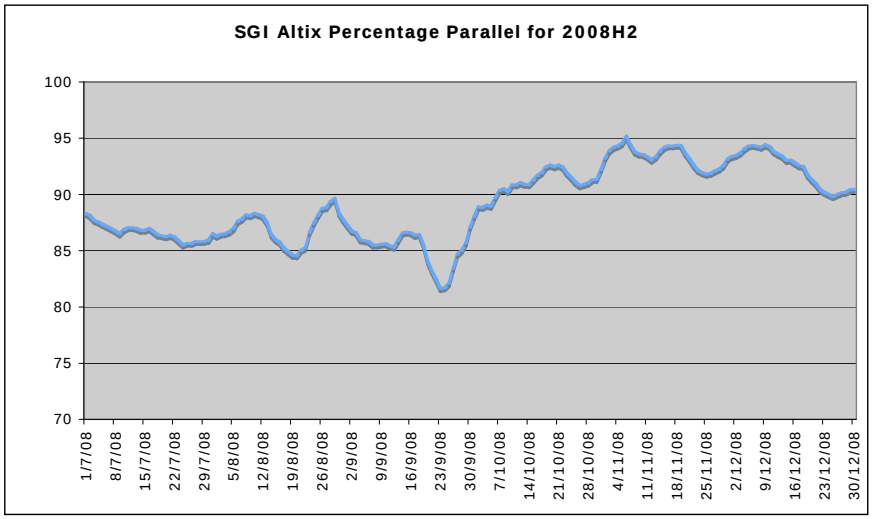

Breakdown of Number of Processors Used on the AC

The following pie-chart shows the fraction of the total usage consumed by jobs making use of the number of processors. The number of CPUs or processors per job increases clockwise from the top, generally in multiples of 2, from 1 CPU / job to 256+ CPUs / job.

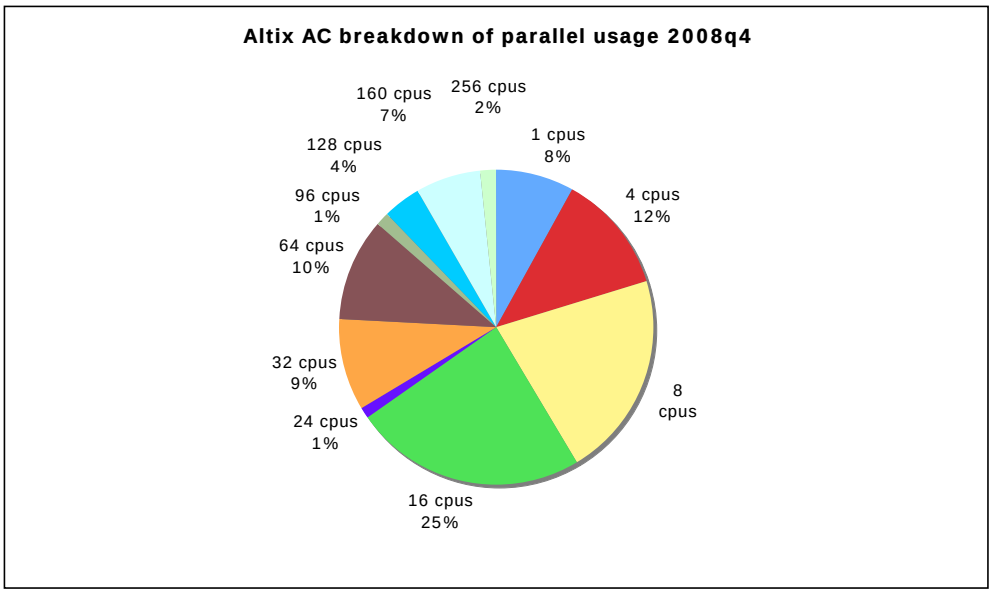

Support and Development with Research Communities and Projects

The following communities have received particular attention in this period.

Biology and Bioinformatics

No special support was provided for this period.

Chemistry

In-depth assistance was provided to Terry Frankcombe (ANU) regarding various theoretical issues with Molpro, Bun Chan (USyd) regarding the GDDI feature in the GAMESS-US program and Neha Gandhi (Curtin) to get her started with the Gaussian package.

Significant prolonged help was spent on various issues with the Molpro 2008 F12 module alerted by Chris Evenhuis (ANU). This was eventually sent back to the Molpro developers in Cardiff and subsequently Stuttgart but is still unresolved.

Quantum chemistry packages installed in this period include updates of molpro, CPMD, ESPResSo, Molden and MRCC on AC; Xe installations of GAMESS-US, NWChem and CPMD and a new installation of OpenBabel on AC.

Jamberoo, previously called the JMolEditor, underwent a number of improvements, which were presented at the eResearch and WATOC conferences. Jamberoo has been down loaded from 45 countries,1022 times July-Sept and 1042 in Oct-Dec. Major work completed in this period included:

1) Initial implementation of output in geom format (Geom format is used for description of polygonal objects and is widely used at Swinburne University of Technology)

2) Implementation of output in VRML format. VRML format files are used for creation of 3d pdf files. In 3d pdf it's possible to embed 3d interactive graphics which can be viewed and manipulated in Adobe Acrobat Reader.

3) Update for Java WebStart automatic download of java3d - To make life of users easier, Java WebStart now will download automatically optional Java components from the SUN website.

4) Implementation of several new atom color schemes

5) Online tutorial on creation of 3d pdf files using Jamberoo (http://sf.anu.edu.au/~vvv900/cct/appl/jmoleditor/manual/tutorials/3d-pdf/index.html)

Other chemistry support:
In-depth support was provided to Terry Frankcombe (ANU: s01) regarding Molpro improvements on Xe, and Damien Carter and Simone Piccinin (MAS:g46), and Alejandro Montoya (MAS:n64) regarding building of several MD program packages including fhi98md, Quantum-Espresso, CPMD.

The Quantum Chemistry package GAMESSPLUS was installed on the AC upon request of Junming Ho. It is based on GAMESS-US

and has a number of extra features, especially in the area of solvation and DFT calculations. There was general support to users from La Trobe, CSIRO, Melbourne, ANU and USyd.

A number of chemistry packages were updated.

- ABINIT 5.4.4 was installed on AC
- Amber-10 was installed on AC,LC,XE
- Accelrys Materials Studio (CASTEP, Discover, DMol3) was updated to 4.3 on AC
- ADF 2008.01 was installed on AC and XE.
- CS-Rosetta was installed on LC
- GAMESS-Plus was installed on the AC
- Gaussview was updated to 4.1.2 on AC, LC
- Gromacs was updated to 3.3.3 on AC
- NW-Chem 5.1 was installed on AC
- Quantum ESPRESSO 4.0.1 was installed on AC

In this period there was also considerable work towards organising and then attending the WATOC2008 conference to be held in Sydney, September 14-19 2008. The National Facility maintained a booth at the WATOC event.

Climate Modelling

He first major project requiring the UK Met Office Unified Model was started. Start-up assistance was provide to Frank Drost (MAS: p14) as he prepares to perform large climate simulations. Assistance was provided for MAS project e14 (England) on several occasions to Shayne McGregor to resolve CCSM3 performance issues on the new Xe system.

Support was provided to MAS project k33 (PI is John You). Support to Nicholas Herold for getting CCSM3 running. They will be using this for paleoclimate modelling so have to generate input files for different scenarios.

Assistance provided to Nicholas Herold (MAS: p42, PI D Muller) to run ancillary programs for generating input data to CCSM3 as part of a Miocene climate modelling project.

Environmental Sciences

There has been ongoing work on the large/continental-scale landscape classification system (Mackey, g42). More information on the research drivers can be found on http://fennerschool-research.anu.edu.au/anuwchub/projects/project01/index.php. The purpose of the work has been to provide landscape researchers with access to a suite of tools and ported to a high performance platform, with such tools scaled to the continental scale. The programming work has primarily take place at ANUSF by Ben Davies with some updates to specific CSIRO codes by John Gallant, with overall steering by Brendan Mackey (ANU) and Ben Evans (ANU) with collaboration from Neil McKenzie (CSIRO).

In this quarter, an improved version of the windows based PCTL tool for ESRI ArcGIS was produced by John Gallant that addresses issues determined in this project. A virtual machine was set up on the data cluster to allow a common environment for users to perform this small step in the classification work.

Further work was completed to automate the data plots that were previously done by hand. A proprietary method of producing the statistics has been replaced with R and this has been rigorously tested to ensure the results are identical up to acceptable tolerances. GDAL has also been used to replace other proprietary functions from the previous system.

As a result of this work, a new continental scale LF7 classification at 9s/250m DEM can now been produced.

Engineering and Fluid Mechanics

Some improvements were made to the parallel performance of Abaqus on the AC, as needed for Hung Kha of ANU project x33. Fluent has been updated to version 6.3.26 on AC and LC.

Consultancy was provided to Andrew Wilkins (MAS m71, PI M Craig) to profile and optimise OpenMP code used for the simulation of underground coal mines.

Plate Tectonics

The latest version of Underworld was ported. This is primarily needed for MAS projects m18 and n69. Underworld is a 3D-parallel geodynamic modelling framework capable of Deriving viscous / viscoplastic thermal, chemical and thermochemical models consistent with tectonic processes, such as mantle convection and lithospheric deformation over long time scales. We resolved various errors in the build process and reported these back to developers.

Consultancy was provided to the pilot Geoscience Australia project (p25) with the porting and optimisation of their tsunami modelling code.

Medicine

There has been ongoing support to Nic Cherbuin ANU:j66 using Freesurfer and FSL to do image analysis and categorisation of MRI brain scans. The code will need to be ported to the Xe and binaries are available for Centos x86-64.

Social Sciences

Mai Pham from ACERH (ANU:p09) was assisted during several runs of GLLAMM using Stata to generate regressions of data on the performance of private and public hospitals. Stata was found to run well on 2 processors allowing her to complete the analysis in a noticeable reduction in time.

General tools and Other packages

Various tools were updated in this quarter - netCDF4.0 (ac), MKL 10.3.020 (ac, xe), Intel Fortran and C compilers version 10.1.018 (ac, xe), NCL (ac), Totalview 8.6.0-1 (ac,xe), GSL1.11 (ac), Matlab R2008a (xe), R 2.8.0 (ac,xe), FFTW 3.2 (ac,xe). There was a close investigation of OpenMPI 1.3 and Vampir for profiling MPI jobs on the XE. Other profilers such as Tau have also been used a comparison.

A range of software on the system is regularly installed and updated as new versions of packages, compilers and libraries become available. The full set of software available can be found on the NCI software registry http://nf. nci.org.au/facilities/software.

System Procurements and Installation

Substantial activity has been undertaken related to two tenders for supercomputers. There was a major tender to replace the peak system at the National Facility and the operational system for the Bureau of Meteorology. This has involved the consideration of the computer room housing the National Facility with regard to power, cooling equipment and external plant, pipework, and transition plans from the space currently used by the Altix 3700 that will need to be cleared for occupancy by the new system.

A second tender was also completed which has resulted in the purchase of a SGI Altix XE 350 system. This system comprises:
- 156 (2 x quad core Intel Harpertown CPUs, 3.0GHz and 1600MHz FSB) nodes each with 16 Gbyte DDR2-800 memory, with some of the nodes with 32 GBytes.
- Total memory is 2.7 Tbytes.
- DDR Infiniband interconnect, which is an upgrade of our existing infiniband hardware
- a Lustre filesystem with 130 Tbytes storage capacity.

The system achieves an aggregate SPECfp_rate_2006 of 12168 over its compute nodes, equivalent to approximately 60% of the AC computational capacity.

The system purchase was completed and successfully completed acceptance tests in December. Early access users, particularly related to CAWCR activities, were then encouraged onto the system. The new Lustre filesystem was being tuned.

System Modifications

SGI System AC – ac.nci.org.au
There has been no major change to the AC in this period.

Dell Linux Cluster LC - lc.nci.org.au
There were no changes to the lc system. However, users are aware that this system will be decommissioned in early January 2009.

SGI Altix XE Linux Cluster
The previous generation of XE system has now been replaced by the new larger system already mentioned. The old Xe hardware was mostly reused for the IO subsystem of the new cluster. There was significant administration and system programming work to prepare the system for production usage. This has mainly involved software to manage the system (using OneSIS), job launching using PBS and OpenMPI and changes to how shared libraries are loaded in a system with multiple versions of compilers. A large number of packages have been ported to the Xe in preparation for production use of the system.

Grid Integration
AC was fully integrated to the grid with the implementation of the Globus Monitoring and Discovery Service, which allows grid users to discover what software and services are available to them from AC. A grid gateway was also implemented for the RSES TerraWulf.

A Shibboleth identity provision service was established to allow users of the National Facility to authenticate to Grisu (http://www.arcs.org.au/GridGrisu) and access the grid.

General Administration Tasks
The transformation of the national facility website was completed and the website is now called http://nf.nci.org.au. This transformation included a new appearance which is harmonised with the main NCI website http://www.nci.org.au. The redesign has been made more manageable by using common headers and style sheets (CSS), avoiding duplication of code, and providing a more flexible way of adding headers and footers to arbitrary pages

The National Facility userguide was also updated to use CSS instead of frames, simplification of page editing and improvement to the navigation.

Operational Data

The AC system was available for 97.36% of the total time theoretically available. A graph of the system uptime is shown below. The theoretical maximum assumes the full system was 100% reliable and that no downtime was required for system management during the entire period.

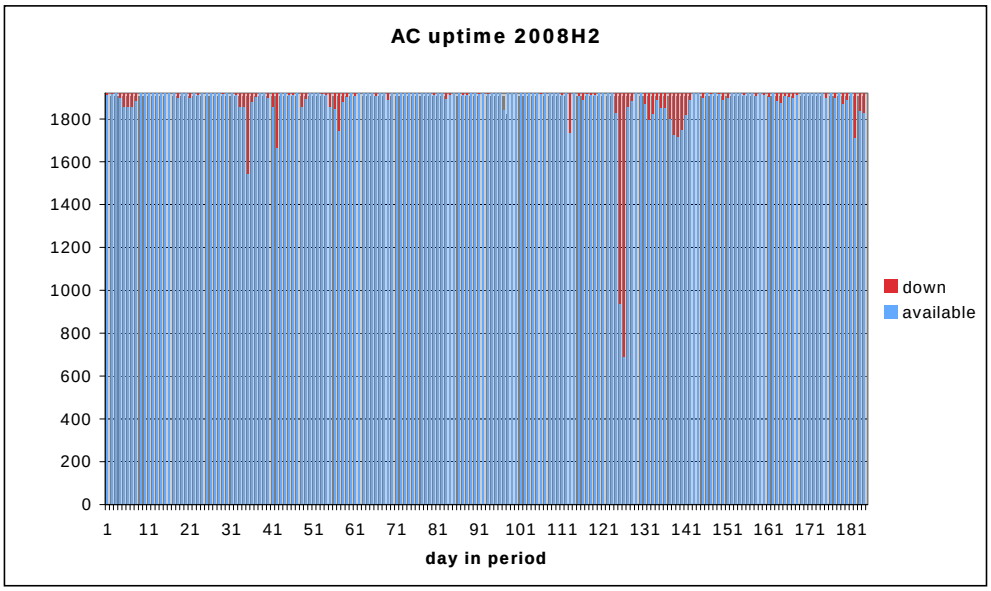

All downtimes are reported on http://nf.nci.org.au/facilities/ac/downtime.php

Datasets hosting and Data Access

In the past the committee had approved MAS-D (data) projects as well as computational storage and important datasets required for computational workflows. Many of these projects are now being reorganized for the new NCI allocation requirements.

Data Allocations

The following storage allocations were made for data storage. The following table and graph shows the requests and use of tape storage space by the MAS-D, MAS-D startup projects, and also the storage required for computational projects from MAS and partner shares. The figures represent the first copy only (two separate copies are kept of all HSM data, so that total amount stored on the system is approximately double that shown).

Storage Tier/Service	Requests(GB)	Granted (GB)	Maximum Usage(GB)
Database storage	28611	28610	8386
Online disk resident*	17480	4874	-
HSM (tape)	649298	547021	405883
Totals	695389	580505	414269

** System programming for reporting on-line data is still being completed.*

The relational databases are typically being used in conjunction with other parts of the storage infrastructure. They may be used for data collection information, metadata housing, fast data methods for indexing into large storage areas, or managing pipelines.

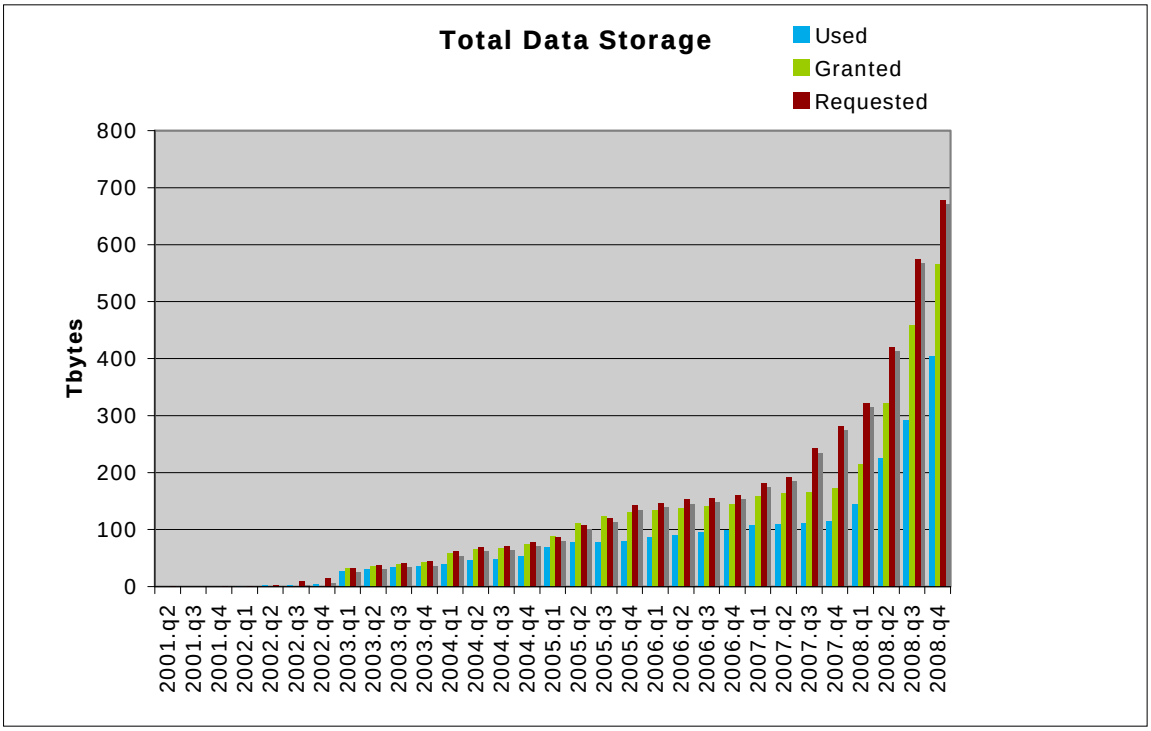

Support for Data

The following communities have received particular attention in this period.

Astronomy

Primary activities in this area revolve around the development of IVOA Standard services to enable access to the MACHO dataset by interfaces and tools conforming to the various IVOA interface standards such as ConeSearch and Simple Image Access Protocol (SIAP). This project work is now well advanced and we expect the previous MACHO web site to be completely functionally replaced by VO compliant versions by early next year. Several key services are now deployed on data cluster, including data stored within a relational database and large data stored on tape.

Work during this quarter focused on performance of the relational database. The largest data product in the MACHO dataset is the photometry data, being time-series data of luminosities for the stars monitored by the project. This dataset consists of data-points number in the hundreds of billions. The initial database design for the VO services involved a single table for this data. As data was ingested into this table, although the query times remained acceptable, the database maintenance operation times became unmanageable. For example the table vacuum operations for managing the storage, consumed well over 100 hours.

This situation has been addressed in several ways. Firstly, the data was moved to a dedicated relational database server. Secondly, the configuration of the relational database server was profiled and tuned. Thirdly, the tables were restructured to allow partitioning of the photometry table into sub-tables, allowing maintenance tasks to be performed on much smaller volumes of data at any one time and data ingest procedures have been rewritten to take advantage of this table partitioning. The data ingest process was also reordered to enable the largest and most commonly queried photometry sub-sets to be brought online first. It is still expected to take several months to ingest the entire photometry dataset.

The existing MACHO web-interface was also maintained and small modifications made. This mostly involved an update of information and messages, fixing some format issues, and reconfiguration to use updated ftp services on the system.

SDSS mirror-node

The data-release six (DR6) dataset has been downloaded, and is being ingested to the data cluster's SDSS mirror service. This dataset replaces all earlier SDSS data-releases.

SkyMapper (MAS p12)
There has been initial development of IVOA Standard services to enable access to the SkyMapper science dataset by interfaces and tools that conform to the various IVOA interface standards such as ConeSearch and Simple Image Access Protocol (SIAP). Development has commenced on a ConeSearch service for the SkyMapper object catalogue and object photometry dataset. These services have not yet been deployed as there is no real SkyMapper data since the telescope is not yet online. The immediate plan is to simulate a reasonable amount of data, in consultation with SkyMapper project staff, and use this data as a demonstration/testing dataset for the VO services as they are developed further.

Environmental Studies
The large MODIS dataset is still being trialled on the data cluster and is used to instigate some interest within CSIRO and Geosciences for computational projects. Geosciences are in the process of transferring data in preparation for several pilot computational projects.

Climate Data
No new changes.

Network Traffic

The following table and graph summaries the data traffic for this period. Nearly all data received was domestic on-net, with data transmitted a mixture of international on-net and domestic on-net in a 1:10 ratio.

Period	Transmitted (Tbytes)	Received (Tbytes)	Total Transacted (Tbytes)
2008q3	21.11	41.41	62.52
2008q4	18.18	11.07	29.25
Total for 2008H2	39.29	52.48	91.76

The maximum data across the network for any one day for this period was 1.8 Tbytes on 29th November. The maximum received and transmitted on any one day was 1.7 Tbytes and 0.57 Tbytes respectively.

The facility has become a national resource for some types of data. Computational output from the AC is mostly transferred to the Data Cluster and some of these datasets are then accessed through specialised data services running on the Data Cluster.

On October 31, the AARNet link to the National Facility was upgraded to a 10Gbit bandwidth connection, the first 10 Gbit link in Australia.

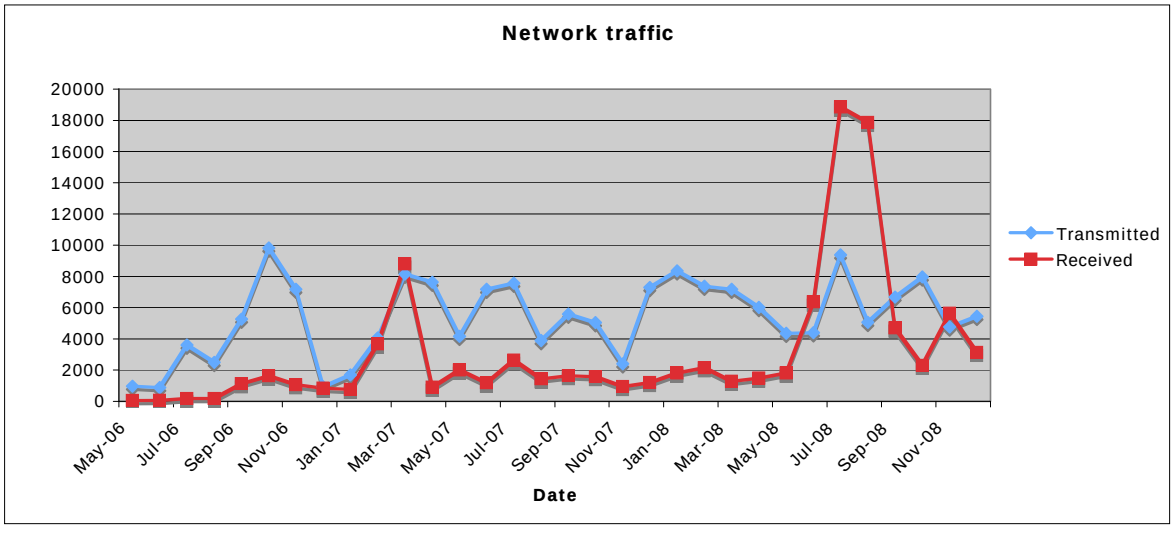

System modifications

Data Cluster – dc.apac.edu.au
The installation issues related to the virtualisation software, VMware ESX, were resolved with a firmware update to the 6140 storage controllers, and hotpatch being generated by VMware as a result of our service call. A windows management host for ESX server deployment has been installed, completing the ESX infrastructure.

The existing D280 disk array is being reconfigured to be dedicated as a metadata device for the DC. Once this is put into production on the DC users will see a noticeable improvement of the metadata performance.

The National Facility is also participating in an early access program for SAM-QFS 5 that is using an implementation of an objects-based filesystem. A test environment was configured and trial load employed. As well as ensuring that the new implementation meets our needs, the test environment will also make sure we are better prepared to make the change once the new filesystem software is ready for production use and support.

The server, previously known as *store*, has been upgraded to Solaris 10 and configured as a node of the data cluster. It is being used for some legacy applications that require a big-endian architecture. In time it will be replaced by a much smaller node.

System Management
No major changes have taken place.

Grid Integration
No major changes have been required.

System Downtimes

The move to a clustered architecture with multiple metadata servers and virtualisation has substantially increased the ability to provide non-stop services for the broad range of data services. We are still improving the system configuration to manage potential system outages. There were two brief downtimes in April, one to resolve a hardware fault was isolated on one of the metadata servers and the other by a non-fenced installation process that reconfigured filesystem partition information. Both issues were resolved, and the former was made easier by migrating the metadata serving to a second metadata server host. The metadata servers do not currently automatically failover. A manual process is currently used to a second hot standby metadata server. We will look to implement automatic failover in the second half of this year, though may delay this to use the new method with SAM-QFS 5.

Downtimes are reported on http://nf.nci.org.au/facilities/mdss/downtime.php

Helpdesk

In the first quarter of 2008 there were 572 emails received by the helpdesk. Around 23% were concerned with minor matters (eg. Passwords, requests for variations to job limits, minor compilation problems etc.) and the remainder being more substantial problems. The median response time was 22 minutes and 25% had a response time greater than four hours (this includes responses arriving on weekends and out-of-hours), but only 4% took longer than 1 day.

In the second quarter there were 595 emails to help, 20% relating the minor matters. The median response time was 20 minutes and 16% had a response time greater than four hours, but only 2% longer than 1 day.

National Facility User Workshops, Training, Education and Visits

Drs Evans, Humble and Singleton attended several meetings to discuss the joint tender with the Bureau of Meteorology. Drs Amos, Kobayashi, Vassiliev, Rostov and Bliznyuk attended the international conference on Computational Chemistry WATOC, held in Sydney. Dr Kobayashi was one of the main organisers of this conference which was well attended by researchers using the National Facility.

Dr Kobayashi organised a Gaussian Workshop on 9-12 Sept and an ADF workshop on 13 Sept, both at Sydney University.

Conference and Meeting Attendance

Who	Where	Why	When
Ben Evans	Sydney	NCI Meeting	02sep
Ben Evans David Singleton	Melbourne	Tender Meeting	08sep-09sep
Roger Amos Rika Kobayashi Vladislav Vassiliev Ivan Rostov Andrey Bliznyuk	Sydney	WATOC 2008	14sep -19sep
Margaret Kahn Stephen McMahon Joseph Antony	Melbourne	eResearch Conference	28sep – 05oct

Jon Smillie Jonathan McCabe Paul Warren			
Ben Evans	Melbourne	Tender Mtg / eResearch Conf	29sep-04oct
David Singleton	Melbourne	eResearch / Training	03-07 Oct 2008
Jason Ozolins		VMWare Technical Meeting	15 Oct 2008
Ben Evans	Melbourne	Bureau (Tender) Meeting	24 Oct 2008
Vladislav Vassiliev		Lecturing at Swinburne University & meeting users	29 – 31 Oct 2008

Courses during July - Dec 2008
9-12 September Gaussian Workshop - University of Sydney (40 attendees) (Dr Kobayashi)
13 September NCI ADF Workshop - University of Sydney (12 attendees)
6-7 October Introducton to the NF, MPI programming – CAWCR, Melb (Dr Singleton)

Visitors to the National Facility during July - Dec 2008
Tim Pugh, Alf Uhlherr, Robin Bowen. – BoM/CSIRO – 3-4 Jul08
Drs Julia Rice and William Swope - IBM Life Sciences (Almaden) – 22 Sept
Dr Vladimir Pelmenschikov – SCM – 24 Sept
Dr Eng Lim Goh – SGI CTO visited department for presentation on 26 November 2008.
Mt Stromlo vacation students (8) toured the facility on 10 December 2008. (Students are 2nd & 3rd year students from various Australian and New Zealand Universities – accompanied by Mt Stromlo Research Fellow, Amanda Karakas.)
National Youth Science Forum visit – ~30 school students toured the facility on 08 January 2009.
National Youth Science Forum visit – ~30 school students toured the facility on 22 January 2009.

Appendix 4: National Facility Operations Report – January 2009 to June 2009

This Appendix outlines the systems and services in the National Facility and the systems and user support provided by ANU staff working in the National Facility during the period January 2009 to June 2009.

Merit and Time Allocations

Merit Allocation Scheme (MAS) for 2009

MAS Committee Updates

Arrangements are being made to prepare for the time allocation process for 2010. This includes the additional new cycles that will become available with the new system.

The Merit Allocation Committee last met in Canberra on November 26, 2008. The Merit Committee Chairman has provided a detailed report to the Steering Committee meeting which is summarised elsewhere in this report. Since the time of the meeting, the Chairman has granted several small increases in time.

Usage Data for 2009H1

Usage in the first quarter of 2009 (2009Q1) was anomalous and unrepresentative of the first half of 2009 since this period corresponded to the commissioning of the new SGI Altix XE system which available for only part of that time. Accordingly, the data reported below is for the second quarter of 2009 (2009Q2) which is regarded as more representative of the usage during 2009.

During that period, the National Facility serviced a total of 228 computational projects through MAS and other shares (including start up grants) involving 977 researchers. There were an additional 65 researchers associated with projects but who did not have a login-names, bringing the total number of researchers involved with projects to 1064. There were 24 new MAS projects awarded time starting in 2009H1 and 10 new start-up projects, one of which has been converted to an MAS project.

The tables below present the overall number of projects and researchers and their grants and usage during the period, broken down by projects under the MAS and through other shares. More details on the individual projects are available on the NF website. It should be noted that some researchers have access to both MAS and another share project. Projects may also have more than one source of grant.

Of these projects, 178 had support from ARC or external grants. (Note that CSIRO projects are not eligible for ARC grants). Of those granted support under the MAS scheme, some 70% of the resources grant were associated with national competitive grants. This is somewhat less that reported last year. Last year's figure was an overestimate, with the number this year should be a more accurate representation based on the better data that has been sought from researchers.

The MAS was 2.22 times over-requested at the start of the period 2008H2. The committee granted 1.2 times the formal MAS share.

Reports on Individual Projects

Principal Investigators of Merit Allocation Scheme projects are required to submit a report on the project each year. These reports are available on-line at http://nf.nci.org.au/annual_reports/. A list of all MAS projects including brief descriptions can also be found at http://nf.nci.org.au/accounts/grants/

Overall Usage of Computational Time

Division	No. of PIs	No. of Projects	No. of Researchers with login-names	No. of Researchers (total)	No. Proj with ARC Support	No. Proj with NHMRC Support	No. Proj with Other Funding	Grant as % of Total Resource Available (*)	Usage as % of Total Resource Used (**)	Formal Percent Share of the System
MAS	151	153	751	799	94	5	45	55.4	60.6	45.95
Other Shares	69	75	358	381	28	1	9	42.5	39.4	54.05
Start-up	10	10	20	21	0	0	0	<0.1	<0.1	<0.1
TOTALS	214	238	999	1064	122	6	54	107.9	100	

(*) All Schemes (MAS and partner Share) are entitled to over subscribe their shares (in a manner comparable to that used by airlines to ensure a high level of seat occupancy). In the case of the MAS, the over subscription factor was 1.2 and all partners are entitled to oversubscribe their shares by this amount. In the first half of 2009, it is evident that partners did not completely allocate their share under the new shares model. This will almost certainly rectify itself in coming periods.

(**) The system scheduler provides a mechanism referred to as "bonus time" that ensures that the system continues to undertake useful work even though some users may have already consumed their allocation. Bonus time is available only if the system is otherwise idle and hence does not impact on access to other users that have resources still available to them. As is evident from the table above, some partner allocations were not completely consumed and so the time notionally allocated to their projects was consumed in bonus time, thus boosting the total share consumed by the MAS to be substantially larger that its allocated share.

Other Shares(MAS included for reference) mas 45.95

Division	No. of PIs	No. of Projects	No. of Users with login-names	No. of Researchers (total)	No. with ARC Support	No. with NHMRC Support	No. with Other Funding	Grant as % of Total Resource Available	Usage as % of Total Resource Used	Formal Percentage Share of the System
MAS									60.6 (*)	45.95
ANU	40	44	232	245	26	1	3	25.6	22.4	26.40
CAWCR	1	1	22	22	0	0	0	3.8	2.5	3.79
CSIRO	21	21	88	92	1	0	5	22.4	13.8	22.17
IVEC	5	5	31	33	0	0	0	0.4	0.5	0.38
QCIF	4	4	7	11	1	0	1	0.4	0.2	0.38

Note: Totals represent unique PI, researchers, and users.

Overall Usage and demand for the System

The following graphs show the amount of queued work on the system in the first half of 2009. There has been approximately 2.6 times the number of CPUs queued on the systems than overall available, rising to a peak of approximately 4 times oversubscribed.

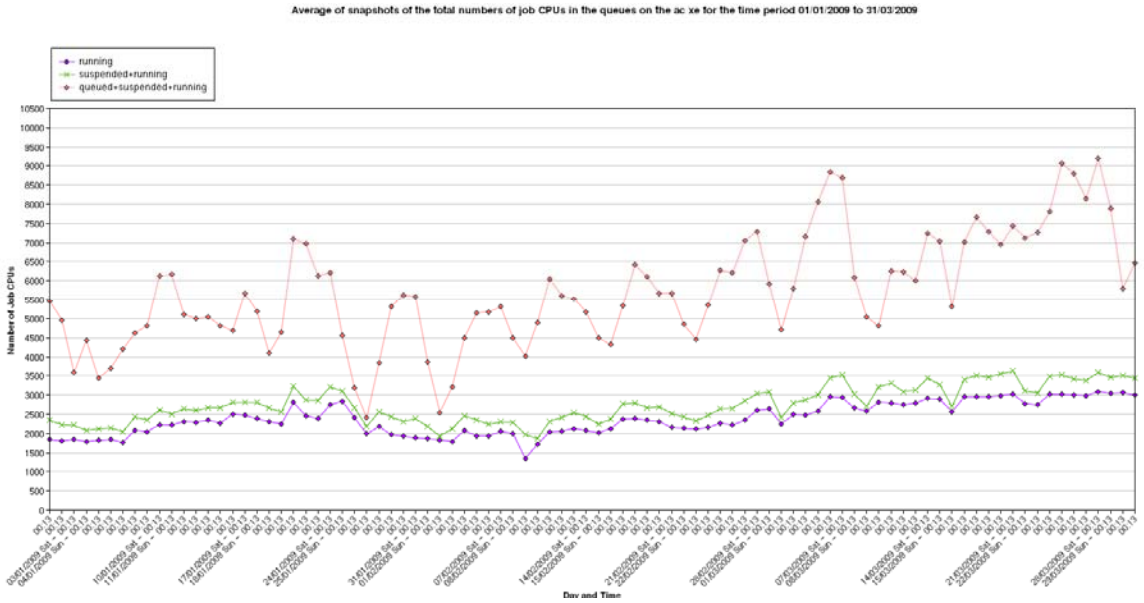

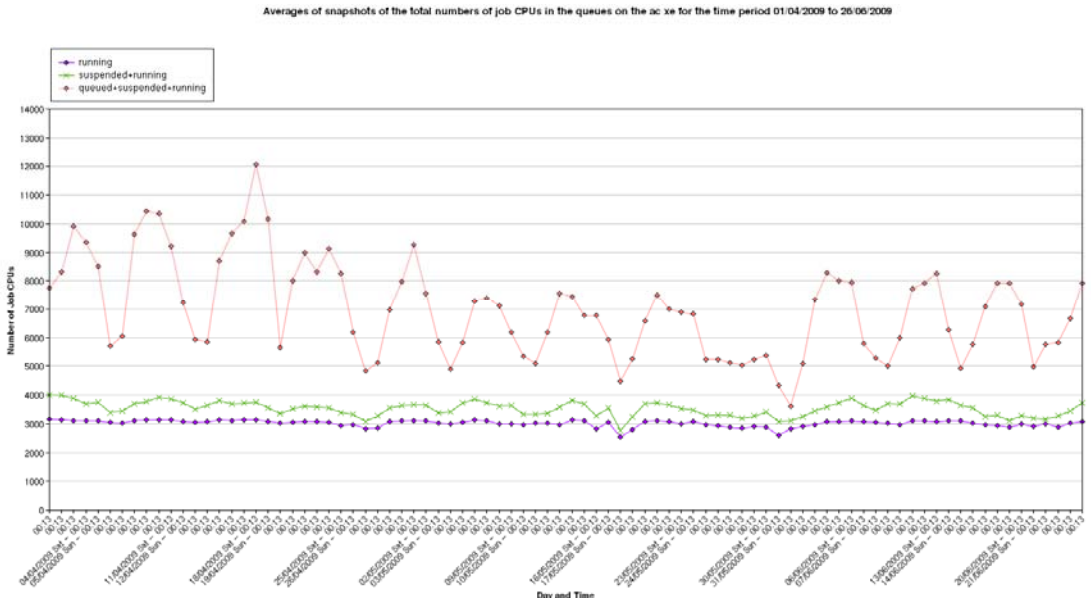

The MAS share has been fully used on the system over the whole period as shown in the following graphs. Overall, system utilisation was an average of 96% utilisation for the period.

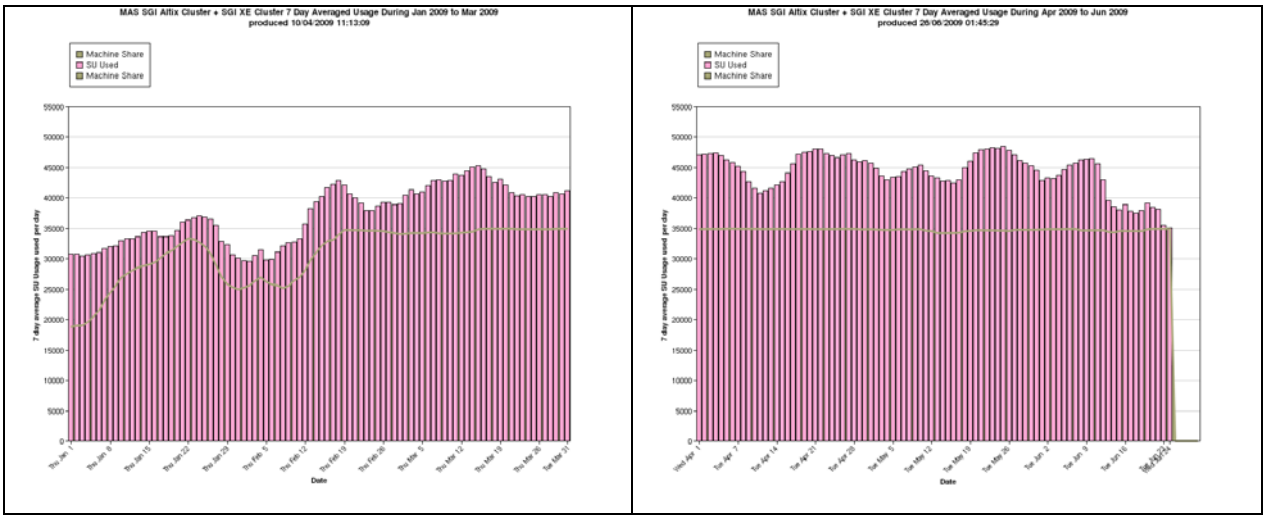

The following graphs show the usage of the National Facility by the various shareholders.

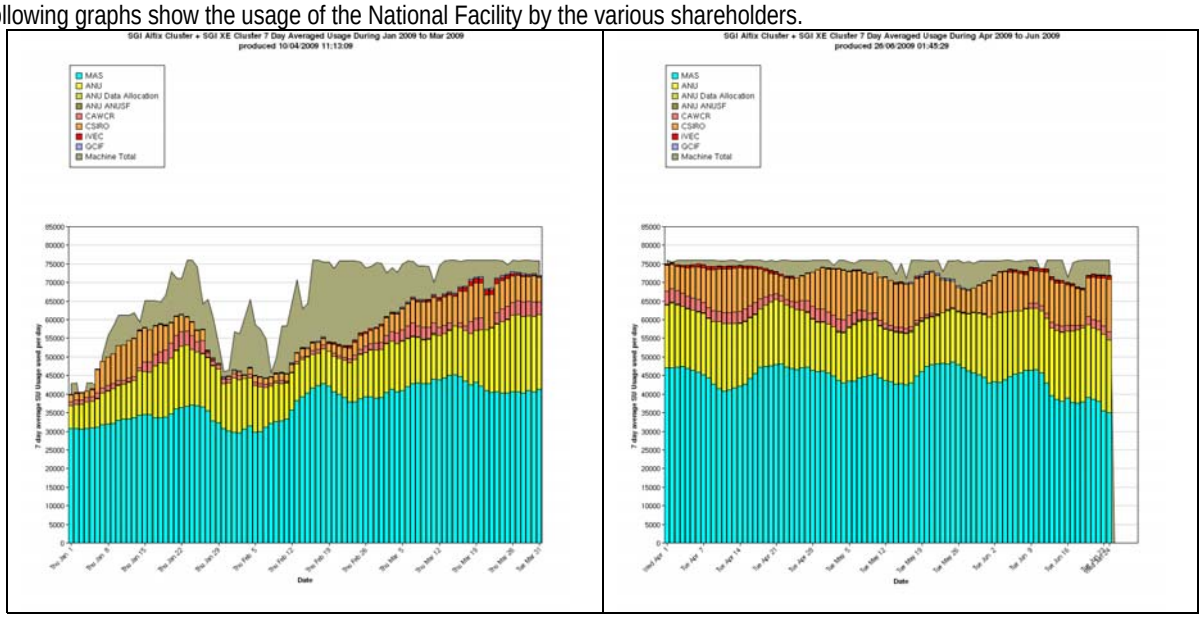

The following graphs show the demand for MAS share over a longer period; 2004 to present. The graph clearly shows the increased demand for peak computing cycles in 2009.

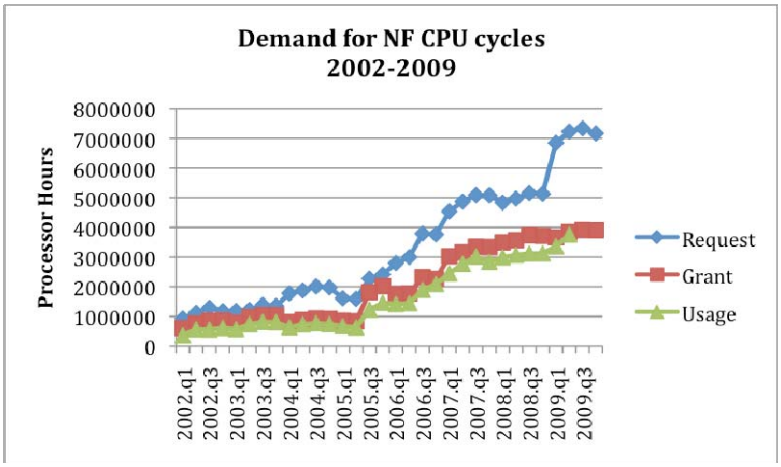

Usage by Research Field

A summary of the usage by ANZSRC Field of Research categories during the period is given in the two pie-charts below. The first set is restricted to MAS computational projects only while the second shows usage over the entire user base.

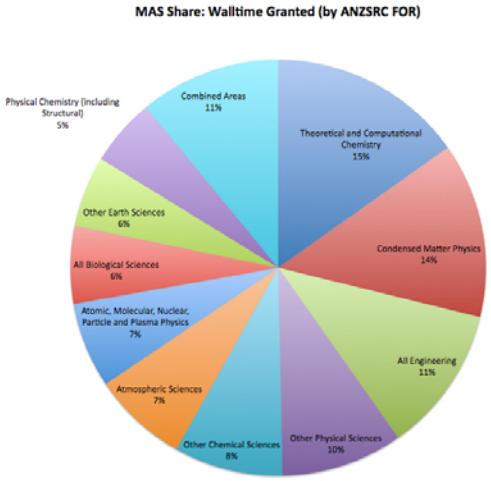

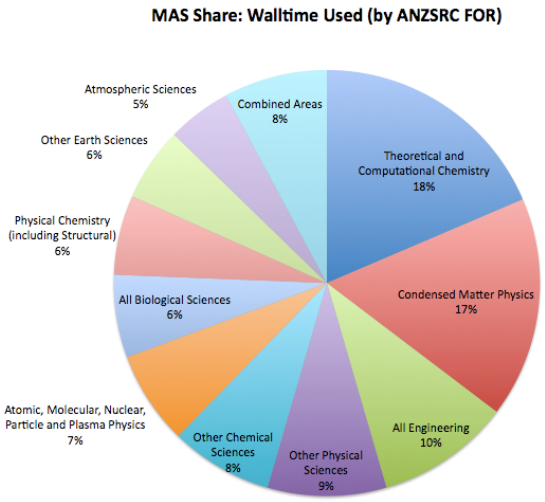

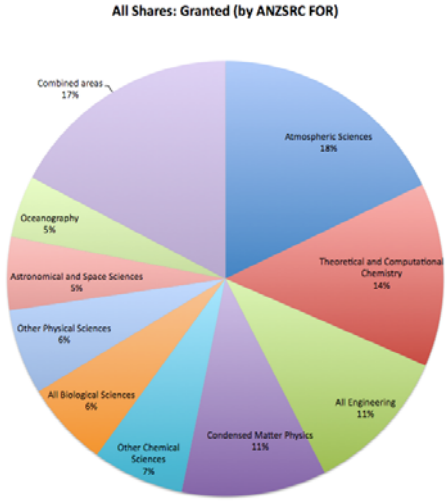

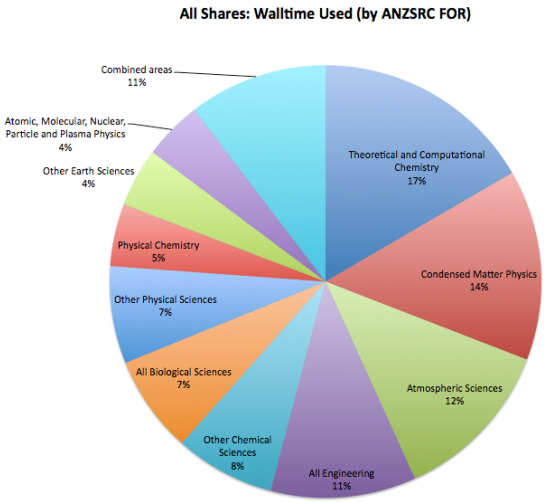

Measure of Parallel Use of the AC

The following graph indicates the percentage of the ac system resources used by parallel jobs. The NF encourages users who might otherwise run an ensemble of tasks as a parallel task, to run the jobs serially to allow the scheduler to make maximum use of the system.

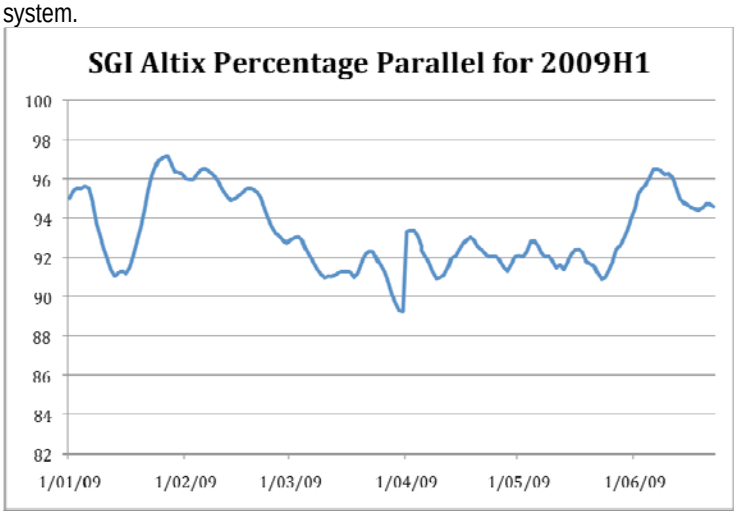

Breakdown of Number of Processors Used on the AC

The following pie-chart shows the fraction of the total usage consumed by jobs making use of the number of processors. The number of CPUs or processors per job increases clockwise from the top, generally in multiples of 2, from 1 CPU / job to 256+ CPUs / job.

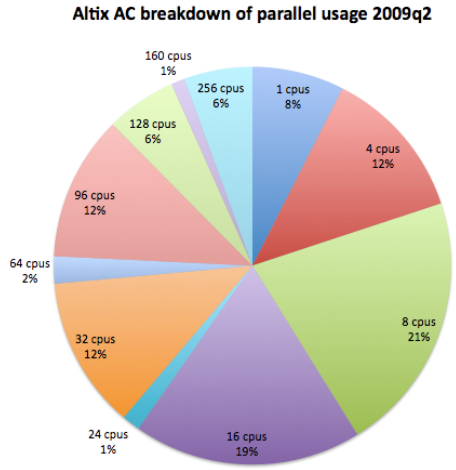

Support and Development with Research Communities and Projects

The following communities have received particular attention in this period.

Astronomy

Work progressed on porting the SkyMapper (ANU:p12) work flow to the XE. This has involved making additions to installed software such as Python and also some changes to the work pipeline code.

The parallel visualization and graphical analysis tool VisIt was installed on the AC for Alexander Wagner (ANU:x49). VisIt can be run remotely from a local workstation with the computationally demanding rendering work being done as a batch job on the AC.

Biology

MAS Project d85, M Ragan. Investigated problems with the 64 bit version of MrBayes on both AC and XE for Karin Kassahn. Installed a patched version of MrBayes that is supposed to handle the 64 bit error in the commands sumt and sump. Tested by Karin Kassahn on large problems but was found to run slower than the non-patched version.

Chemistry

In-depth help was provided to Terry Frankcombe (ANU:s01) regarding various issues with the performance of Molpro through customized builds. Aspects such as the filesystem from which Molpro is run and the impact of other user jobs on the nodes were investigated.

Special support was provided for users from various institutions (Monash, USyd, La Trobe, Swinburne, ANU, CSIRO, UWA, RMIT, UNSW, Curtin). Extensive support was given to Matthew Stewart of La Trobe University to identify problems with the ACES_MAB program, and to Wenyi Yan of Monash University to help wirh ABAQUS problems on Xe.

A customised build of Gaussian03 was prepared for Jeff Reimers, USyd.

Ivan Rostov from ANUSF exposed some possible thread safety issues when linking Gaussian g03e01 built with the Intel compilers to the latest Intel MKL libraries.

In-depth support was also given to Simone Piccinin (MAS:g46) regarding building of the cp2k quantum chemistry package, to David Poger (MAS:m39) on building Gromacs and to Brad Wells (MAS:n29) on porting a code for molecular Monte Carlo adsorption simulations. Andrey Bliznyuk from ANUSF provided information on building Vasp, an ab-initio simulation package, to several users.

The software package DFTB+, version 1.0.1 was compiled and installed on both ac and xe following request of Amanda Barnard (asb595, CSIRO). The DFTB+ is fast and efficient stand-alone implementation of the Density Functional based Tight Binding (DFTB) method.

The software package Dylax, version 5.2.7 was compiled on both ac and xe following request of Amanda Barnard (asb595, CSIRO). Dylax is an earlier implementation of the Density Functional based Tight Binding (DFTB) method, which currently replaced by DFTB+. However, according to Amanda Barnard's request, Dylax has some unique features missed in DFTB+. Since the Dylax program is considered now as obsolete, the created Dylax binaries were copied to the user's directory on ac and xe, rather than installing them in the central location of the supercomputers' filesystems.

Some work on the ANUSF chemistry product “Jamberoo” was completed:
1) Implementation of parsing of Amber PREP files
2) Implementation of parsing of Amber Lib (Off) files
3) Tutorial on how to visualize volumetric data from the ADF 2008 calculations
4) Development of JavaScript code for creating of the dynamic 3d pdf content (see attached pdf sample with dynamic content)

ANUSF staff member, Dr Vlad Vassiliev, submitted a joint paper with Feng Wang. Dr Vassiliev manuscript "Towards Interactive 3D Graphics in Chemistry Publications" was also submitted for publication.

The software package VMD, version 1.8.6, was installed on XE. VMD is a molecular visualization program for displaying, animating, and analyzing large biomolecular systems using 3-D graphics and built-in scripting.

A number of chemistry packages were updated.
- ADF 2008.01 was set to the default on AC and Xe.
- Amber 9 and 10 were installed on the Xe
- An optimized version of Gaussian g03e01 built with the Intel compilers was installed on the Xe.
- Dalton was installed on the Xe.
- The default version of Siesta was updated.
- ABINIT 5.6.5 on Xe
- LAMMPS on Xe
- Molden 4.7 on Xe

The contract for Turbomole was re-established with the purchase of Turbomole 6.0 and utility program TMolex for Xe.

Climate Modelling

There has been some in depth analysis of the climate codes that have been installed on the system for CAWCR. The codes have been ported to the Xe system without problems and now produce correct results. Martin Dix reported a compiler bug that was then fixed in the latest release of the compilers and updated on the XE. Several perl modules were also installed to allow the Unified Model to work.

Continuing support was provided to Frank Drost (MAS: p14) as he sets up and starts to run large climate simulations using the UK Met Office Unified Model.

Help was given to Stephen Jeffrey (MAS:n71) to debug scripts written to submit a sequence of jobs.

Advice was given to Michael Bates (MAS:e14) on porting the Princeton ocean modeling code POM to the XE.

Various versions of the NCAR Command Language (NCL) were installed on AC, AC-X86 and XE for Jiafu Mao (CSIRO:n89,MAS:p78)

for graphical pre and post processing of NetCDF files from climate simulations.

Earth System: Plate Techtonics

Underworld 1.2.0 was installed on both AC and XE for Monash/VPAC and also to be used over the grid. The new build process initially failed on the AC but succeeded after changes were made to the Scons source code in order to allow the Scons configure process to execute on the AC. This correction was relayed to John Mansour (m18 MAS a subproject of MAS p67) and John Spencer at VPAC. The earlier version Underworld 1.1.0 was reinstalled so that it could be used over the grid.

Environmental Sciences

Ecological Regionalisation of Australia / Continental Classification System (ANU:g42): Much of the work for Q2 was carried out in conjunction with FSES research assistants who were working on a second stage of the 'Bush Heritage' project. New users were provided with direct support on using the ANUSF systems for their research. This included writing a major update to the user guide thus providing a useful reference for non-technical users.

Further work involved evolution of the LF7 terrain analysis/classification method. This allowed for production of a finer grained continental terrain layer to represent the Australian landscape more accurately. This layer will be used directly in upcoming landscape classification work. Other development includes reworking some system outputs in order to produce additional statistical outputs. These form the basis of tools to assist in exploring research topics in the bio-ecoregionalisation field.

Engineering and Fluid Mechanics

Support was provided for users of Abaqus on the AC and XE, in particular to n42 (MAS, PI is Wenyi Yan), p37 (MAS, PI is Brian Falzon) , x33 (ANU, PI is S Kalyanasundaram). Significant support was also given to users who incorporated user-defined functions into their Abaqus jobs to ensure that they worked.

An intermittent failure of 8 CPU jobs was investigated and resolved on the XE. Abaqus runs twice as fast on the XE as on the AC but 8 CPUs is the maximum number that can be used. Several of the engineering packages, Abaqus, LSDyna and Fluent, can be difficult to run efficiently and experience in optimal use of these packages is steadily being developed.

OpenFOAM was installed on the XE for Project d77 (MAS Hugh Blackburn). The associated graphical viewer is not currently operational.

The computational fluid dynamics code CFD++ has been installed on the XE and is to be tested by users from Russell Boyce's hypersonic research group (MAS:n23). Investigated the module MB_CNS as part of this work. Sequential version built on the AC but there were problems with the parallel version. User instructions on how to build on the XE were provided. Ultimately we referred the user to the developers who had previously been an active group on the AC (Peter Jacobs, g60). Various problems with managing the license for CFD++ were also resolved.

The most recent version of the finite element code LS-Dyna was installed on the XE.

Earth Sciences

Assistance was given to Andy Hogg (MAS:g40, ANU:x77) in getting MPI code running on the XE and the performance tuned. We also helped Marshall Ward in the same project to build and run the application GOLD, a specialized geophysical fluid dynamics model. This project was also supported by the development of data manipulation tools for the MIT GCM model (for atmosphere, ocean, and climate studies). Pre-existing tools were investigated and found to be incapable of performing the necessary concatenation of data tiles. A custom tool was developed to support the research, allowing RSES staff to undertake further research.

Nicholas Jones (MAS:m19) was given substantial support in techniques for handling and storing the large data sets required to run the atmospheric chemistry code GeosChem.

The 1.1.0 version of Underworld, a 3-D parallel framework for modeling plate tectonics, was reinstalled on both the AC and the XE to suit the requirements of project (MAS:m18).

Grads, a free software package for plotting in earth sciences was installed on the XE and AC-x86 for Xiaobing Zhou (CAWCR:p73).

Economics

Project ba0 (Manfred Lenzen is the PI). This project is currently in start-up phase to investigate MPI optimization code.

Medicine

There has been ongoing low level support to Nic Cherbuin (ANU:j66) in using Freesurfer and FSL to do image analysis and categorisation of MRI brain scans on the XE. As a result of this it was determined that the XE is over twice as fast as the LC and all scripts are now running successfully. Advice was also given on handling his large data sets.

Research in Optimisation

Help has also been given to Ting Yu on MPI parallelism and using numerical maths libraries. Ultimately this should be a large parallel optimization problem but it is still some way off.

General tools and Other packages

Various tools were updated in this quarter - MKL 10.1.1.019 (XE), Intel Fortran and C compilers version 11.1.0.083 (XE), NCL 5.1.1 (AC), FFTW 3.2 (XE),Scons, PetSC 2.3.3-p12 with C++ support (XE),TAO 1.9 (XE), NETCDF 3.6.2 (XE), GCC 4.3.3 (XE).

The MPI profiler IPM, Integrated Performance Monitoring, is being tested as a light weight profiler for MPI jobs on the XE. It has been successfully used for a CCSM3 climate code run. The next step is to modify IPM so it can handle times when the MPI job is suspended by the queue scheduler.

Assistance was given to Michael Dudalev (ANU:n97) in using Mathematica for large jobs. There are some outstanding problems caused by occasional suspension of Mathematica licenses from running jobs on the XE. The most recent version of Mathematica was installed on the XE and some work done to investigate its parallel behaviour.

SunStudio12 and SunStudio Express were installed on the XE in order to investigate the performance of the Sun compilers.

FFTW3.2 was rebuilt with shared libraries on the XE for Mark Abraham (MAS:d52).

A range of software on the system is regularly installed and updated as new versions of packages, compilers and libraries become available. The full set of software available can be found on the NCI software registry http://nf.nci.org.au/facilities/software.

System Procurements and Installation (Status Report as of July 1, 2009)

Substantial activity has been undertaken related to acquisition of the new Sun supercomputers. The project will not be reported here in detail. However, the basic plan is to install the system in multiple phases:

Stage 1a: Two new Sun compute racks with fully populated blades will be delivered, which provide more computational power than the existing SGI. On the current implementation timelines the system will be operational in September.

Transition and removal of the existing SGI: Following acceptance of the Sun Stage 1a, the user data and user information will be moved from the SGI to the Sun. Details on this transition will be mailed to the users to ease the process for all in this major hardware replacement. The transition for users should be eased as users have already had access to the SGI Xe system which is essentially the same as the new computer, processors and with the ANU software stack and filesystem.

Due to the need to replace the cooling system, the SGI will be totally removed and the space cleared for the new cooling pipework to be installed. A new false floor tile will also be installed during this period, which has a much higher weight loading. The new tiles are also needed as there are not enough old tiles that can be cut for the new system with the strength required.

The existing cooling system and a modified pipe will be used to allow Stage 1a to continue operating. There is a risk that the heat generated by both Stage 1a and the Xe during this period will require a reduction in service. Plans to mitigate this have been scheduled (including the modified pipe, and extending the existing Carrier cooling plant maintenance) and further action may be required as the situation is monitored. The Xe may be moved within the machine room at a later date to facilitate further floor tile replacement, and this will result in a downtime some time in 2010.

Stage 1b: The second, and larger, component of the new Sun system will then be installed and connected to the new pipework connected to the new cooling plant. Power upgrades have been taking place and are well within schedule. The new facility cooling plant facility construction and component delivery is currently on schedule and proceeding within expectations. Once connected, the new system will undergo installation and acceptance tests. It is planned to be ready for full use in January 2010. The transition of users from Stage 1a to 1b will be relatively minor, but will require an acceptable downtime. The Stage 1a system will be merged into the final system as soon as possible once the system is stable.

System Modifications

SGI System AC – ac.nci.org.au

There has been no major change to the AC in this period.

SGI Altix XE Linux Cluster – xe.nci.org.au

The Xe is a first generation of cluster without local disk attached to each node and no general solution currently exists from a vendor. A considerable amount of system design has therefore been needed to ensure a working model was reached.

The design of the OS layout Xe system was designed by the NF staff and includes a novel use of both a root-on-Lustre Operating System, as well as components that a resident within a ramdisk. The current system is based on CentOS 5.3.

Before the system was released the system team spent time diagnosing faulty filesystem server stability issues and disk subsystem issues. Some rigorous test regimes were implemented to help in the future. SGI sent replacement parts for a number of faulty disk subsystem components.

The operating system model that has been developed has formed the basis of the design to be used on the blade-based peak Sun supercomputer. The Xe system has been invaluable in our preparation for that system, as well as providing a critical increase in the cycles available to users. The development has provided feedback that has assisted our understanding and production experience used in discussions with Sun for the final peak system design – including MDT (metadata) and OSS (storage) performance. Several software stack improvements were also made to help prepare for the Sun system. The new system will use the ANU system stack which is now well suited to the needs of the National Facility users as used in production.

We acknowledge the assistance/perseverance of some of the early test users of the Xe system as we ironed out the system bugs. The climate community (and particularly CAWCR) were able to make significant use of the machine during the period before it was made available for general use.

The Xe system was released to general users on March 4. Since then the system has been remarkably stable (over 99% available), and provides a contrast to the difficulties experienced with some of the proprietary components of previous systems. The transition has been very smooth for users, especially those from the old linux cluster LC.

A special filesystem sharing node was connected to the Lustre filesystem to help support a visualisation render workflow to the ANUSF Vizlab.

Dell Linux Cluster – lc.apac.edu.au
The LC system was switched off on 31 March. Users were moved across to the Xe system along with their data. The system was deconstructed and removed.

Grid Integration
No major upgrades have been necessary. There have been some changes to allow new groups to access the NF service.

Operational Data

The AC system was available for 98.16% in 2009.q1 and 99.61% in 2009.q2 of the total time theoretically available (overall 98.89%), and the Xe system was available 72.18% in 2009.q1 and 99.13% in 2009.q2 (overall 85.66%). Graphs of each system uptime is shown below. The theoretical maximum assumes the full system was 100% reliable and that no downtime was required for system management during the entire period.

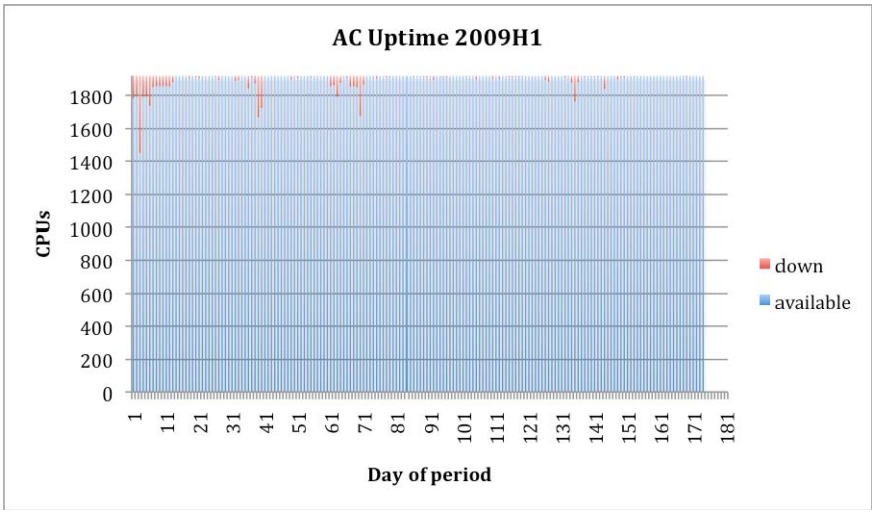

All downtimes are reported on http://nf.nci.org.au/facilities/ac/downtime.php

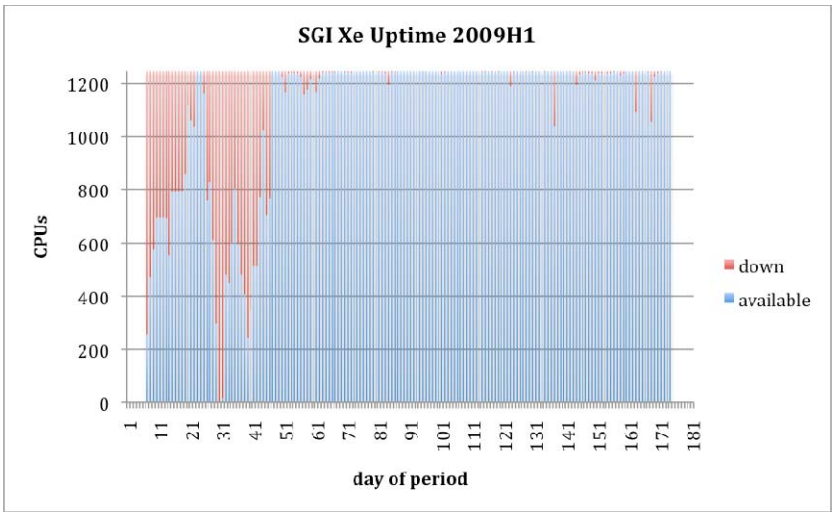

During January and February, the Xe system, and particularly the Lustre filesystem and operating system implementation was still being tested. The system was released to general use on March 4. Since then has achieved high levels of availability, similar to the AC.
All downtimes are reported on http://nf.nci.org.au/facilities/xe/downtime.php

Datasets hosting and Data Access

Data Allocations

The following storage allocations were made for data storage. The following table and graph shows the requests and use of tape storage by the various schemes for computational projects from MAS and partner shares and dataset hosting. The figures represent the first copy only (two separate copies are kept of all HSM data, so that total amount stored on the system is approximately double that shown).

Storage Tier	Requests(GB)	Granted (GB)	Maximum Usage(GB)
Database storage	29437	28736	15457
Online disk resident*	5696	5447	-
HSM (tape)	663556	550107	338407
Totals	698689	584290	353864

The system programming for reporting online data is still being completed.

The relational databases are typically being used in conjunction with other parts of the storage infrastructure. They may be used for data collection information, metadata housing, fast data methods for indexing into large storage areas, or managing pipelines.

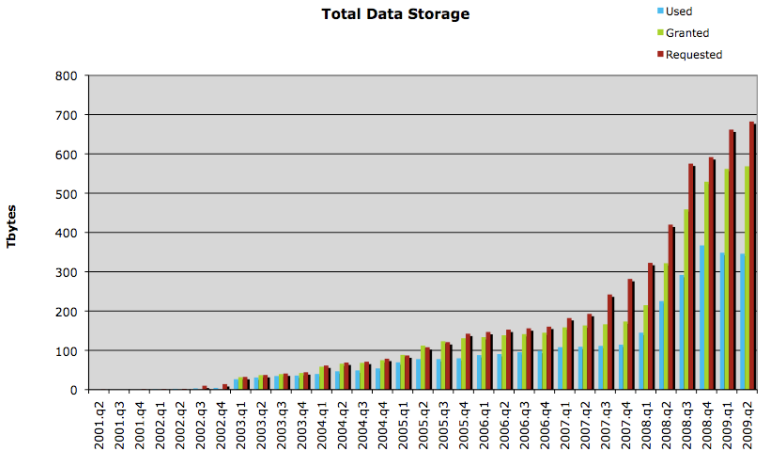

Support for Data

The following communities have received particular attention in this period.

Astronomy

The previous MACHO web site has now been functionally replaced by a VO compliant version. Several key services are now deployed on data cluster, including data stored within a relational database and large data stored on tape. The remaining development is to provide a new data ingest procedure for the MACHO photometry data. Close to half of the total data is now available via the MACHO VO photometry service.

Now that the MACHO VO services are in a production release state, the expected associated management and maintenance work on these active services now constitutes am increasing percentage of the workload related to these services.

SkyMapper (MAS p12)

With the SkyMapper project approaching first-light, there has been significant activity to complete the configuration of the SkyMapper Data Analysis Pipeline (SDAP) software, and the final implementation of data transfer service between Siding Springs and ANUSF.

SDAP configuration has progressed on several fronts. Firstly, the extensive list of software dependencies has been installed on the xe system. Secondly, the SDAP software itself has been installed, and where necessary, the software has been compiled, including third-party packages. Finally, integration testing is progressing, with a standard set of test data and test cases being run to enable detailed configuration of the pipeline to be performed.

Systems Development

Systems

ANUSF Data Cloud – dc

Activities this period have largely focused on:
(a) the consolidation of services and servers for both hosting and database environments.
(b) the creation of new services at the IP network layer.

The system now hosts web services using ESX server, which in turn relies on the cluster file system for its data resources. Database servers are also being tested within a VMWare environment to aid in system consolidation.

All production data services such as GridFTP, SRB and the legacy FTP are now running on native QFS clients.

A DDN 6620 array that arrived in January has had a number of issues. We have analysed its limitations and it will now be used a (fast) SAM disk archive copy. By the end of the year we plan to consolidate our existing Dell server setup into a blade server chassis for database and DC cluster work.

An additional disk space will be installed over the next quarter to meet ongoing demand for the data cloud service.

The system deployment model has been developed to rolling out the services with unique addresses that can be moved across system nodes without interruption to users.

System Management

VMWare management is now undertaken using a centralized host within our management subnet.

ARCS Grid Integration

Globus was upgraded, including the GridFTP server, and implemented on dgal1. It has been tested and verified to be working.

Data Services

Data Transfers

The commands netmv and netcp have been installed on the Xe to set up batch jobs for transferring large data sets between this system and the data cloud. Additional integrity checking is being added to the mdss commands for data copying to and from the NF data cloud. This will provide users will alerts as to possible corruption of data.

A technology for high-speed data transfers between the ANU, CSIRO and the Bureau is being finalised with the plan to have in place for production services on the new systems.

LDR

The LDR service used for international LIGO was reconfigured within the Data Cloud to provide better system integration. The IP system changes have meant that ldr.dc.nci.org.au is now functioning. The LDR service is also used to transfer MODIS data from Perth onto the DC as one of the environmental datasets required, in particular for CSIRO users.

The international data transfers to JPL for Solar FTS data (m19) has been improved. The transfers of geos-chem data for the m19 project have benefitted.

FTP
The DC now provides the ability for anonymous ftp "drop boxes", which is available on ftp-dc.nci.org.au.

SRB
The SRB service was removed from service as users have moved on to better technologies on the system.

OpenDAP
The operational OpenDAP server that provides mostly IPCC Climate data will be shifted to a dedicated node soon due to increased load. The OpenDAP server is planned to be linked into other international partners as part of the service required for the IPCC international data federation, replacing the older centralised PCMDI site. This federation comprises the sites producing the large model output which need to be shared for further analysis.

Ruby on Rails
The Ruby on Rails environment is being packaged for easy redeployment and movement through the data cloud.

System Downtime

There have been some interruptions to the service in this period. These were due to underlying hardware instability in the disk array arising from faulty firmware. Sun escalated the resolution, which required a new version to be sent. However, the new firmware took some time to arrive and resulted in several unplanned outages in January and March.

Downtimes are reported on http://nf.nci.org.au/facilities/dc/downtime.php

Network Traffic

The following table and graph summaries the data traffic across the AARNet network for this period. Nearly all data transferred was on-net.

Period	Transmitted (Tbytes)	Received (Tbytes)	Total Transacted (Tbytes)
2009q1	30.64	11.63	42.27
2009q2	49.35	13.36	62.71
Total for 2009H1	79.99	24.99	104.98

The maximum data across the AARNet network for any one day for this period was 6.8 Tbytes on 7th May. The maximum received and transmitted on any one day was 0.9 Tbytes and 6.4 Tbytes respectively.

The data cloud has become a national resource for some types of data well beyond that required for the computational output of the local computational system.

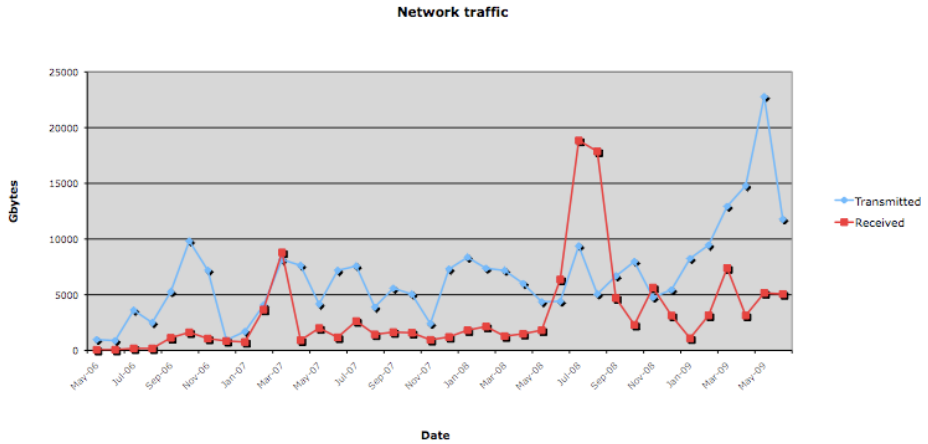

Note, the statistics do not show the data transferred between the computational systems and the local data cloud. Also, on October 31, 2008, the AARNet link to the National Facility was upgraded to a 10Gbit bandwidth connection, the first 10 Gbit link in Australia.

National Facility Helpdesk

In the first quarter of 2009 there were 675 emails received by the helpdesk. Around 21% were concerned with minor matters (eg. Passwords, requests for variations to job limits, minor compilation problems etc.) and the remainder being more substantial problems. The median response time was 21 minutes and 21% had a response time greater than four hours (this includes responses arriving on weekends and out-of-hours), but only 6% took longer than 1 day.

In the second quarter there were 570 emails to help, 20% relating the minor matters. The median response time was 23 minutes and 18% had a response time greater than four hours, but only 5% longer than 1 day.

National Facility User Workshops, Training, Education and Visits

Conference and Meeting Attendance

Who	Where	Why	When
Ben Evans	Sydney	Supporting the Data Lifecycle symposium	11 February 2009
Kevin Pulo	Hobart	Presentation at Australian Linux Conference 2009.	18-30 February 2009
Ben Evans Margaret Kahn Andrey Bliznyuk Rika Kobayashi Ben Davies Joseph Antony David Singleton Jon Smillie Roger Amos Ajay Limaye Drew Whitehouse Vladislav Vassiliev Ivan Rostov Olaf Delgado-Friedrichs	Rydges Lakeside Hotel, Canberra	CSIRO Computational Workshop	17-18 March 2009
Ajay Limaye	Monash Uni	Drishti Workshop	27 March 2009
Ben Evans	Sydney	NCI Tender brief and bureau discussions	27 March 2009
Robert Davy Joseph Antony Jason Ozolins	Sydney	Dell Meeting – research server purchases	29 April 2009
Ben Evans	Amsterdam, & UK	DEISA / PRACE Symposium, Hector Discussions, Visiting Computer Centres	07-29 May 2009
Rika Kobayashi	Helsinki, Oslo & UK	ICQC Conference, Satellite Workshop & Visiting Computer Centres	16 Jun-12 Jul 2009

Courses during January - June 2009

Drishti Workshop, Monash University – x19 attendees on 27 March 2009.

Visitors to the National Facility during January - June 2009

National Youth Science Forum visit – ~30 school students toured the facility on 08 January 2009.
National Youth Science Forum visit – ~30 school students toured the facility on 22 January 2009.
02 March 2009 – Salvy Russo
04 Jun 2009 – Marc Hamilton, Vice President of Americas Systems Practice in Sun's Global Sales and Service organization.