



NATIONAL COMPUTATIONAL INFRASTRUCTURE

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

# **ANNUAL BUSINESS PLAN**

**Combining NCRIS and EIF Project Plans**

## **2010–11**

National Computational Infrastructure  
CSIT Building, North Road  
Australian National University  
Ph: +61 2 6125 9800  
[www.nci.org.au](http://www.nci.org.au)

31 March 2009



## TABLE OF CONTENTS

1	Executive Summary .....	1
2	Introduction.....	5
3	Status of the Projects .....	6
3.1	NCRIS NCI Project.....	6
3.1.1	National Facility Program.....	6
3.1.2	Specialised Facilities Program.....	7
3.1.3	Computational Tools and Techniques Program .....	7
3.1.4	Outreach .....	8
3.1.5	Partnership.....	9
3.1.6	Planning, Access and Policy.....	9
3.2	EIF Climate HPC Project.....	9
4	Governance and Management.....	11
4.1	Governance of NCI.....	11
4.2	Management of NCI .....	11
4.3	EIF Project Governance .....	11
4.3.1	Project Steering Committee.....	11
4.3.2	Planning Groups .....	12
5	Activities to Progress the NCRIS and EIF Projects.....	15
5.1	Peak Systems — Planning, Acquisition and Commissioning (NCRIS/EIF) .....	15
5.1.1	General Purpose GPU Trial (NCI CT&T, NCRIS) .....	15
5.1.2	Expansion of the Current Peak System (EIF).....	15
5.1.3	Storage Infrastructure (EIF) .....	15
5.1.4	EIF Climate HPC System Infrastructure and Services Planning .....	16
5.2	Data Centre — Planning and Construction (EIF).....	17
5.3	Specialised Facilities (NCRIS) .....	18
5.3.1	Background.....	18
5.3.2	Present Status .....	18
5.3.3	Services: Specialised Facility in Bioinformatics .....	18
5.3.4	Services: Specialised Facility in Imaging and Visualisation .....	19
5.4	Specialised Support and Development.....	19
5.4.1	Earth Systems Science (NCRIS CT&T).....	20
5.4.2	Astronomy (NCRIS CT&T).....	21
5.4.3	Cloud Computing Services (NCRIS CT&T) .....	22
5.4.4	Code Scaling Project (ANU/NCI-Oracle-Collaborative Agreement) .....	23
5.5	Access (NCRIS and EIF) .....	23
5.5.1	Background and Context .....	23
5.5.2	Access Schemes .....	25
5.5.3	Allocations for Calendar Year 2010 .....	26
5.5.4	National Merit Allocation Scheme.....	27
5.6	Outreach Activities .....	28
5.6.1	Public Relations and Promotional Materials (NCRIS) .....	28
5.6.2	Building and Maintaining Relationships (NCRIS) .....	28
5.6.3	Building Capability (ANU/NCI-Oracle Collaborative Agreement) .....	29
5.6.4	Conferences (NCRIS).....	30
5.7	Planning and Policy Development .....	30
5.7.1	Industry and Commercial Access Policy.....	30
5.7.2	Redevelopment of the Partner Service Agreement .....	30
5.7.3	Resource Allocation Model .....	30
6	Project Resources .....	32
6.1	Partnership and Collaboration.....	32
6.1.1	Current Status.....	32
6.1.2	Developing the Co-investment Base .....	32
6.1.3	Industry/Commercial Income .....	33

6.2	Staffing Projections .....	33
6.3	Financial Projections .....	35
6.3.1	NCRIS NCI Project Budget .....	35
6.3.2	EIF Climate HPC Budget .....	40
7	Milestones .....	41
7.1	NCRIS NCI Project .....	41
7.2	EIF Climate HPC Project .....	42
8	Risk Management .....	44
8.1	NCRIS NCI Project .....	44
8.1.1	Adequacy of Resources .....	44
8.1.2	Services .....	44
8.1.3	Security .....	45
8.1.4	NCRIS Objectives .....	45
8.2	EIF Climate HPC Project .....	46
8.2.1	Financial — Infrastructure and Sustainability .....	46
8.2.2	Suitability for Purpose .....	48
8.2.3	Construction .....	49
8.2.4	System Procurement .....	51
8.2.5	Services .....	52
9	Key Performance Indicators .....	54
9.1	Performance Indicators — NCRIS NCI Project .....	54
9.2	Performance Indicators — EIF Climate HPC Project .....	56



## 1 Executive Summary

This coming period, the 2010–11 financial year, is a transition year for NCI, corresponding to the conclusion of funding under the NCRIS NCI Project, and the commencement of the EIF Climate HPC Project. Accordingly, this document integrates both activities within a single framework—the NCI Business Plan—and seeks to carefully distinguish the respective funding streams, activities, milestones, performance indicators and risks.

The NCRIS NCI Project is now quite mature, with all aspects of its program either now in place, in progress, or completed. In contrast, at the time of writing, the EIF HPC project is about to formally commence, following the signing of the Funding Agreement by the Commonwealth of Australia and the Australian National University (ANU).

In what follows, the key activities of each program that are planned for the coming year are summarised under the relevant headings.

### Project Governance (EIF)

The NCI Steering Committee has recommended to ANU a project governance structure for the EIF Climate HPC Project comprising:

- A Project Steering Committee comprising executive representation from the major stakeholder organisations (ANU, CSIRO and the Bureau of Meteorology), and the chairs of the project management groups, which have carriage of individual aspects of the project.
- Project management, which will be undertaken through the work of three planning groups which report to the Project Steering Committee. These are in the areas of (a) Services and Infrastructure Planning, (b) Data Centre Planning and Implementation, and (c) Systems Planning and Implementation. Valuable work has already been undertaken by the Services and Infrastructure Planning Group which has identified and characterised key science challenges and national benefit drivers.

### Peak Computational Infrastructure (NCRIS and EIF)

Key aspects of the development of the peak systems infrastructure at NCI are:

- Commissioning of *vayu* (Sun Constellation system)—initially installed and commissioned with one-eighth of its final capacity in September 2009, augmented to one quarter of its final capacity in early February 2010, and again to three-quarters of its capability in early March 2010, and finally reaching its full capability (with a peak performance of 140 Tflops) in mid-April 2010. The facility has been warmly received by the research community which is impressed with the performance that it provides.
- Planned performance upgrades in 2010—with the capability of *vayu* being augmented to around 200 Tflops later this year through funding under the EIF contract. Also to be provided for through EIF funds will be upgraded data storage infrastructure (disk and tape), particularly to satisfy the burgeoning needs of the Earth System Grid node.
- Preliminary scoping of the nature, capability, environmental requirements, size and cost of the petascale infrastructure, to be commissioned in 2012, has already been undertaken through expert analysis within NCI (also involving discussions with vendors and international high-end installations).

The coming year sees a number of major milestones including the development of:

- comprehensive services and infrastructure plan (in consultation with experts from the earth systems science and other research communities),
- Collaboration Agreement by 2010Q4 to provide for the recurrent operations of the centre,
- requests for proposals (RFP) for the petascale system infrastructure in 2010Q4, and the public release and subsequent evaluation of the RFP in 2011Q1 and 2011Q2 respectively.



### **EIF Climate HPC Data Centre (EIF)**

Significant progress has already been made in planning for the data centre through the engagement of external consultants to identify site options, and to scope and analyse the data centre requirements based on the anticipated performance and environmental characteristics of the petascale system, as well as the available funding envelope.

ANU will be in a position to provide the Commonwealth with its recommendation for the site for the Data Centre by mid-May 2010, and will use the remainder of 2010 to plan, design, and acquire the necessary approvals for construction. A twelve month construction period covering the year 2011 is anticipated, in preparation for the fitting out of the data centre, and the installation and commissioning of the petascale system in the first half of 2012.

### **Specialised Facilities**

Contracts were signed in late December 2009 with the University of Queensland (and QCIF) and Monash University, the respective proponents of two Specialised Facilities in:

- Bioinformatics (involving UQ, QCIF, CSIRO and NCI),
- Imaging and Visualisation (involving Monash University, Australian Synchrotron, CSIRO, VPAC, NCI).

The necessary infrastructure is being procured for each facility and, under the respective contracts, the services in each of the Specialised Facilities will commence in the latter half of 2010.

NCI is investing \$1.2M in each of the Specialised Facilities and its share of the services from each will be made available to the research community through the NCI Merit Allocation Scheme, for which arrangements are already underway.

### **Computational Tools and Techniques / Specialised Support**

The CT&T program will undergo a change of name to Specialised Support in order to more accurately reflect its role, particularly in the future (EIF funded) environment in which developmental activities of this type are of paramount importance.

The CT&T funded support programs in earth system science (ESS) and astronomy will be fully underway in 2010Q2, with all staff in place — two at the National Facility (i.e., one associated with each program), and two in Melbourne, one at Monash University for the ESS program, and one at Swinburne University for astronomy and astrophysics.

A vital addition to the National Facility Support Team will be provided through the ANU/NCI-Oracle Collaborative Agreement (associated with the contract between ANU and Sun/Oracle for the peak system). This position, from mid-2010, will focus on developmental activities, the thrust of which will be to enhance the scalability of existing codes, to improve access to high scaling codes that can exploit the new infrastructure that is being put in place, and also to increase the uptake in high throughput computing for data intensive applications.

The need for sufficient resources in future years to fund a development team of a size sufficient to derive the maximum benefit from the planned petascale infrastructure is emphasised in the Business Plan, and is budgeted for from 2011 onwards. This initiative will proceed, however, only after sufficient funds to maintain the development team have been put in place under the Collaborative Agreement.

### **Access**

A critical facet of NCI's planning for the coming year will be the development of principles to provide for access to serve the EIF priority area of climate change, earth systems science, and national water management. This together with the implementation of the flagship allocation scheme will substantially change the access landscape.

With the development of multiple petascale systems in Australia within two years, a National Merit Allocation Scheme is an emerging priority in order to provide for coordinated, convenient access for



the Australian research community. NCI expresses its willingness to assist in facilitating this, working with DIISR and the operators of other national high-end systems.

### **Outreach**

With assistance from funds provided through the ANU/NCI-Oracle Collaborative agreement, NCI will seek to foster capability in interdisciplinary research that develops the uptake of high-end computing. This will be done through focussed thematic development programs that directly benefit the stakeholder community of NCI in areas such as the earth systems sciences, geosciences, the environment and possibly others.

NCI will also seek to increase its strategic relationships with high-impact research communities and collaborations (e.g., ARC Centres of Excellence, other national eResearch providers, etc.), and the research organisations that support these. Such a strategy aims to increase the uptake and impact of high-end computing in high profile research in Australia, to more closely embed NCI services in the practice of high-impact research, and, in doing so, to garner additional resources with which to develop the services offered by NCI.

The upgrading of the public face of NCI will continue in the coming year, with a further round of “research highlights” to be developed, and with further efforts to gain media coverage for the research supported by NCI.

### **Planning and Policy Development**

Key aspects of policy development in the coming year include:

- Development of a comprehensive policy for industry and commercial access, and particularly the formalisation of industry/commercial partnerships. NCI anticipates the recruitment of its first industry partner during the coming year.
- Redevelopment of the Partnership Agreement and the Resource Allocation Model, consequent upon the Collaborative Agreement needed to sustain the recurrent operations of the petascale system, and also the introduction of new principles underpinning allocations to the priority area.

### **Collaboration and Partnership**

The development of the partnership that sustains NCI will be a critical task for 2010–11. The primary focus will be on the development of a Collaborative Agreement that will underpin the recurrent operations and the development of NCI in the petascale environment from 2012. This agreement is required as an expenditure milestone (2010Q4) of the EIF Climate HPC Project, permitting the commencement of the Data Centre construction. Informal discussions with the key stakeholders (ANU, CSIRO and the Bureau of Meteorology) indicate a willingness to provide for the necessary co-investment.

In the shorter term, Intersect Australia (of NSW) will join as a partner from 2010Q2 (for a four year period), underpinned by a repurposed ARC LIEF grant and annual subscriptions from the consortium. Geoscience Australia is also anticipated as a new partner during 2010–11.

### **Risk Analysis**

This Business Plan presents a comprehensive risk analysis of the two programs.

#### NCRIS NCI Project

While there are few residual risks associated with the NCRIS NCI Project until its formal completion date (30 June 2011), there are significant risks associated with the continuation of the operation of the (current) peak system at the National Facility beyond that time, compounded by the termination of the current Partner Service Agreement at the end of 2011. Renewal of the current Partner Service Agreement will be considered during negotiations that lead to the Collaborative Agreement through which the recurrent costs of the EIF funded infrastructure are to be provided.

More problematic, however, is how NCI will provide for the ongoing operations of the merit-based access schemes—the Flagship Scheme and the Open Access Scheme—beyond that time. While



partner access to the system will be secured through a Collaboration Agreement and a renewed Partner Service Agreement, the merit-based access schemes can only be provided for through a share that is allocated on the basis of (some fraction of) the Commonwealth's investment determined by a revised Resource Allocation Model to be developed during the course of 2010–11.

NCI is conscious of the need manage the risk associated with the provision of computational support for high impact projects supported by these schemes (80% of which are supported by project funds from the National Competitive Grant Schemes—predominantly ARC) in the post-NCRIS environment. In this regard, the absence of an ongoing source of (Commonwealth) funds, other than the modest residual (or fraction thereof) at the termination of the NCRIS agreement, with which to provide or supplement merit access is of concern.

Accordingly, NCI signals to DIISR the need to discuss these matters during the coming year.

#### EIF Climate HPC Project

The major risks that have to be mitigated in the EIF infrastructure project are:

- meeting the recurrent operational costs of the EIF Climate HPC Project,
- increasing the expert support devoted to developmental activities (software development activities, code scaling, advanced data management etc).

The latter is emphasised throughout this plan, since without such a team, it is unlikely that the potential benefits of access to petascale computational infrastructure can be realised in full.

Provision is made within the NCI Budget for the establishment of such a team commencing in 2011. This will proceed, however, only following the negotiation of the Collaborative Agreement.



## 2 Introduction

National Computational Infrastructure (NCI) was established in June 2007 under a Funding Agreement between the Commonwealth of Australia and the Australian National University. With funding provided initially through the NCRIS Platforms for Collaboration (PfC) Investment Plan, and subsequently supplemented through the substantial investment of a number of co-investing partners, NCI provides state-of-the-art, high-end computing facilities and services to the Australian research community, building on the successful Australian Partnership for Advanced Computing (APAC) National Facility (NF) program.

Its plans and operations are informed by its mission statement:

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

The Commonwealth Budget, announced on 12 May 2009, included a package of three Super Science Initiatives as part of the Government's response to the Review of the National Innovation System. These investments, formulated from consideration of the "Strategic Roadmap for Australian Research Infrastructure", were developed during 2008 by the Department of Innovation, Industry, Science and Research (DIISR) following consultation with Australia's research sector. In responding to the identified need for additional high performance computing capacity in targeted fields, the Super Science (Marine and Climate) initiative incorporated a \$50M allocation over three financial years 2009–12 for "Upgrading the Climate High Performance Computing Facilities". This was described in the Budget papers as follows:

"Australia's peak High Performance Computing Capacity will be upgraded. This new high end computing infrastructure will be used by the Bureau of Meteorology, CSIRO, ANU and others to address climate change, earth systems science, and national water management issues. A next generation building and machine room will be required to house the upgraded system."

Subsequently, the Commonwealth Department of Innovation, Industry, Science and Research (DIISR) approached ANU to negotiate an Education Investment Fund (EIF) agreement to establish a new data centre and a petascale high-end computing facility. In its formal reply to DIISR (in a letter dated 20 November 2009), ANU noted that "the University plans to pursue contract objectives through an extension of the role of the National Computational Infrastructure ... by expanding the present model that underpins NCI to provide for a broader range of climate and marine stakeholders".

The financial year 2010-11 spans the final year of the NCRIS NCI Project and the first year of the EIF Climate HPC Project. Since both projects operate under the governance of the NCI Steering Committee, and since there are many elements in common, a single NCI Business Plan has been written, with this document providing the Business Plans for both the:

- NCRIS NCI Project, and
- EIF Climate HPC Centre Project.

Its structure combines the requirements of each individual plan, and clearly distinguishes the funding lines where appropriate.

Separate budgets for each project are provided.



## 3 Status of the Projects

### 3.1 NCRIS NCI Project

The past year has seen substantial progress in each aspect of the NCRIS NCI project. This is summarised under the relevant headings below.

#### 3.1.1 National Facility Program

The planned upgrades to the National Facility peak system are nearing completion. These commenced in February 2008 with the decision by ANU and the Bureau of Meteorology (BoM) to proceed with a joint tender for interoperable systems to provide for the operational needs of BoM, and for a digital laboratory to accommodate advanced computational research in the earth systems sciences, and other fields of research which have a requirement for high-end computational services.

A public tender was let for the joint systems in April 2008, and closed in late May 2008, with a detailed evaluation phase continuing until late September 2008. Contract negotiations commenced with Sun Microsystems in October 2008 and continued through until early March 2009 at which point the contract was signed. This period coincided with the deteriorating international financial environment and necessitated substantial risk mitigation measures to secure the transactions, the bulk of them having exposure to US currency. The combination of the need to secure currency transactions at a known fixed price, and the subsequent improvements in the value Australian dollar against the US dollar has led to currency exchange losses, reported previously in the 2008–09 Annual Report.

The initial delivery of hardware for Phase 1 (one-eighth of the final systems, and sufficient to take over the computational load of the former peak system with a margin of 50%) occurred in late July 2009. By mid-August 2009, the five rack system (2 C48 processor racks, 1 M9 Infiniband rack, and 2 racks of disk storage) had been assembled and cabled, and had passed diagnostic tests. Staff from the National Facility was provided with access to the system from mid-August to build and test the software stack, to configure the Lustre filesystem, and to test live user codes.

With the stability and acceptance tests (comprising 120 hours of continuous running) successfully completed, ANU formally accepted the system on 16 September 2009, after which all users were transitioned to the new system within the space of two working days. Despite the brevity of the transition period, there was no inconvenience to the user community since the SGI Altix XE, acquired earlier in the year as a development / transition system, provided a functionally equivalent platform on which to test and redevelop codes in readiness for the new environment. The transition to the new facility proved to be a very smooth experience for researchers who commented favourably on the strong performance of the new system and also on the straightforward and transparent migration process. Representative comments to this effect were included in the NCI Annual Report for 2008–09.

From late September 2009 onwards, there was substantial work undertaken to refurbish the Leonard Huxley Data Centre (new power and cooling services, and a computer room floor), and to commission the new plant (associated building, water chillers and new power infrastructure).

The hardware for Phase 2 of the installation was delivered in mid-October 2009, with the system assembly and cabling being completed by mid-November 2009. Originally, it had been anticipated that the system would pass acceptance tests by late December 2009. During December, however, in Sun's testing of the system components, problems in the functionality and reliability of the interconnect fabric were identified. Initially, these were suspected to be related to inappropriate handling of the optical fibre cables (connecting compute nodes to the Infiniband switch), but by January 2010 it was clear that there were systemic problems with the connectors on the optical fibre cables. During January 2010, extensive investigations by Sun Engineering and its optical fibre component supplier definitively identified the source of the problem which lay in the active end-connectors on the optical fibres, some of which were significantly outside of their nominal impedance specification.



Throughout this process, NCI was in contact with its major stakeholders, whose production loads had been delayed, and also communicated with the wider user community via broadcast email.

In the week commencing 8 February 2010, a field remediation procedure which reprogrammed the firmware in the connectors was undertaken, after which NCI and Sun staff configured a revised first stage system with twice the previous computational capacity. This system was made available to the research community on 15 February 2010 in order to alleviate the backlogs and bottlenecks which had arisen.

Sun/Oracle has since remediated the faulty switching fabric in its entirety, and has commissioned the whole of the second phase of the installation, which successfully passed its acceptance tests on 9 March 2010. The system was formally accepted on 10 March 2009 after which, over a period of three days, the filesystem was transferred from the smaller (Phase 1) to the larger (Phase 2) system. The service to the research community (with a system of three-quarters of its final capacity) resumed on the evening of 13 March 2010. It is anticipated that the complete (140 TFlop) system will be available from mid-April 2010 once the Phase 1 and Phase 2 installations are “merged”.

A survey of users of the new National Facility peak system is planned for 2010Q3.

### 3.1.2 Specialised Facilities Program

Contracts for the Specialised Facilities in Bioinformatics and in Imaging and Visualisation were negotiated during the latter half of 2009, and executed during the final working week of December 2009.

NCI is investing \$1.2M cash in each of the specialised facilities, for services to be taken over a period of three years (2010H2–2013H2). The partner organisations in each of the two specialised facilities are:

- Bioinformatics (involving UQ, QCIF, CSIRO and NCI),
- Imaging and Visualisation (involving Monash University, Australian Synchrotron, VPAC, CSIRO and NCI),

with the total investment in each, and the agreed services to be provided, being encoded in the Schedules of the respective contracts. A summary of this information, particularly the services and resources to be provided, is listed in the tables in Secs 5.3.3 and 5.3.4 respectively.

The University of Queensland has just executed a contract with SGI for the provision of compute and storage infrastructure upgrades, while Monash University will shortly release a request for proposals to provide for its infrastructure upgrades.

The services in each of the Specialised Facilities are scheduled to become available from mid-2010, although this may be delayed by two months due to delays in provisioning the infrastructure. In preparation for this, access arrangements for implementation through the NCI Merit Allocation Scheme are being put in place.

### 3.1.3 Computational Tools and Techniques Program

In consultation with the Climate Science / Earth Systems Science (ESS) and Astronomy communities, project plans for each of the ESS and Astronomy support projects have been developed.

The ESS project (see Sec. 5.4.1) is now in place, providing for two full-time staff for two years—one located at the National Facility, and another in Melbourne (the hub of the development of the ACCESS system) to be hosted by Monash University. A staffing agreement to provide for the Melbourne based position was executed with Monash University in the final business week of 2009 and an appointment has been made already. The National Facility position is in the process of being filled.

The Astronomy project (see Sec. 5.4.2) will be operational in 2010Q2. Again, there are two full-time positions for two years—one within the NCI National Facility and the other at the Swinburne Centre



for Astrophysics and Supercomputing—with the agreement of Astronomy Australia, with which the project plan has been evolved. At the time of writing, the staffing agreement to provide for the Swinburne position is being considered by the Contracts Office of Swinburne University.

The Cloud Computing project (see Sec. 5.4.3), underway since last year, is investigating new computational research services to be offered by NCI in high-throughput, data intensive computing, potentially through commercial cloud computing providers. The areas on which effort is presently being focused include astronomy, earth systems science, and bioinformatics, with 1.5 EFT provided for through the CT&T budget, and with up to \$75K per annum available to meet associated expenses (software, access to commercial cloud services etc).

### 3.1.4 Outreach

The past year has seen a substantial upgrade in NCI's public face, with a new image, a new website, and a suite of new publicity materials being developed. The centrepiece of the new publicity materials comprises the research highlights — vignettes of high-impact research undertaken on the NCI National Facility, and written up by a professional science journalist in order to be widely accessible and media friendly. This suite of materials will be further developed in the coming year, as will upgrades to the website.

A highlight of 2009 was the public launch of the new peak computing infrastructure by the Minister for Innovation, Industry, Science and Research, Senator, the Hon. Kim Carr, at a ceremony on 16 November 2009. Other speakers at the occasion included Emeritus Professor Mark Wainwright (Chair, NCI Steering Committee), Dr Alex Zelinsky (representing the CSIRO CEO), Professor Robin Stanton (representing the ANU Vice-Chancellor), and the NCI Director, Professor Lindsay Botten. The speeches emphasised the importance of advanced computational infrastructure in addressing problems of critical importance to the human kind, its role in the advancement of technology and the economies of technologically advanced nations, ANU's strong and sustained stewardship of APAC and NCI, NCI's strong engagement with, and strong support of, the nation's research communities, the strong and growing partnership of ANU, CSIRO, the Bureau of Meteorology, and the future development of a comprehensive digital laboratory for climate change / earth systems science and national water management through the new Super Science initiative.

Following the official proceedings, the Minister had the opportunity to meet and speak with three high profile researchers whose research is supported by NCI:

- Dr Amanda Barnard, CSIRO (Nanotechnology) — winner of the 2009 Malcolm Macintosh Prize for Physical Scientist of the Year,
- Professor Brian Schmidt, ANU (Astronomy) — current Federation Fellow and winner of the 2000 Malcolm Macintosh Prize,
- Professor Andy Pitman, UNSW (Climate Science) — Co-Director, UNSW Climate Change Research Centre,

and also to inspect the new facilities.

There was strong media coverage of the event and of the new facility, with a large number of articles appearing in the press (leading daily newspapers, other press outlets and also Business Review Weekly), and with a number of interviews on ABC Radio (including the PM program on Radio National) being conducted on the day of the launch.

The current year has also seen development in NCI's international relationships, with very productive visits undertaken to the University of Edinburgh (EPCC/Hector, e-Science Institute) and the Engineering and Physical Sciences Research Council of the UK.

Another significant feature of this year's outreach activities has been the pilot programs conducted jointly with Geoscience Australia (GA). As a consequence of these, GA will join NCI as a partner in the coming year.



### 3.1.5 Partnership

The year 2009–10 has seen significant growth in the NCI partnership. The most significant event has been Intersect's decision to join as a partner, a decision made by the Intersect Board in May 2009, following the consortium's success in attracting ARC LIEF funding for new HPC infrastructure. The past year has seen the repurposing of the LIEF grant (so as to acquire a share of the NCI National Facility), with the consent of the ARC, the development and signing of an ARC Multi-Institutional Agreement by all parties (i.e., the original university partners of the grant, and also ANU as host of NCI), and Intersect's signing of the NCI Partner Service Agreement (anticipated in April 2010).

Monash University also joined as an NCI partner during the course of the 2009-10 financial year.

### 3.1.6 Planning, Access and Policy

Amongst the significant activities in planning, policy and access considered by the Steering Committee during the course of the past year were:

- The development of a commercial access framework, including a code of conduct for NCI partners for commercial use of their share,
- A revised access framework to provide for flagship allocations that support high impact research, and which also seeks to upscale the computational sophistication and usage of the peak facility.

## 3.2 EIF Climate HPC Project

This is the first Business Plan to provide for activities under the EIF Climate HPC Project. At the time of its preparation, the Commonwealth Government and ANU are on the point of signing the Funding Agreement.

While there is no formal progress that is to be reported, substantial work has been already undertaken in preparation for a formal start of the Project in the second quarter of 2010. In addition to the planning and project governance framework that is reported in this Business Plan, there is significant preliminary work already underway that is detailed under the following headings and which is also included in the EIF Climate HPC Project Plan and Initial Implementation Plan.

### Background

The project was formally initiated through a letter dated 7 September 2009 from DIISR to ANU inviting the University to formally negotiate an EIF contract, and also seeking its response on a number of matters. In its formal reply to DIISR, in a letter dated 20 November 2009, ANU noted that "the University plans to pursue contract objectives through an extension of the role of the National Computational Infrastructure (NCI) ... by expanding the present model that underpins NCI to provide for a broader range of climate and marine stakeholders".

Since then a contract has been negotiated between ANU/NCI and the Commonwealth. In addition, a Project Plan, and an initial Implementation Plan have been developed, and approved by the NCI Steering Committee at its 9 March 2009 meeting.

### Consolidation of *vayu* (Sun Constellation)

Plans have been developed to consolidate the present peak system (which has just been commissioned) and to increase its capacity by around 25-30%. Plans have also been developed to investigate the role that general purpose GPU technology may play in services to be offered by NCI in coming years.

NCI is further planning to augment the data storage available to the facility, particularly to accommodate the burgeoning needs of the Earth System Grid, through additional disk capacity and new tape robots to be acquired as part of the EIF project.

Scoping of the capacity has already been undertaken and indicative quotations obtained.



### **Services and Infrastructure Planning**

A Services and Infrastructure Planning Group, comprising strong representation from the climate change and earth system science community, was established informally in 2009 to scope the science and national benefit challenges that the petascale infrastructure would be required to meet. This group has already undertaken valuable work which will continue immediately following the signing of the contract.

In addition, a high level specification of the petascale infrastructure that might be acquired in 2012 within the available funding envelope has been scoped in order to inform the environmental requirements of the new Data Centre.

### **Data Centre Planning**

Since the announcement of the Super Science (Marine and Climate) Initiative in the May 2009 Commonwealth Budget, ANU has commissioned two data centre scoping projects, one in late 2009 to examine possible green field sites for the data centre, and a second, more recent study to identify possible sites on the ANU campus.

The latter study, commissioned in March 2010, will scope the building design and logistics within the available funding envelope and will inform ANU's choice of the site, to be communicated to the Commonwealth by the mid-May 2010 milestone date.



## 4 Governance and Management

### 4.1 Governance of NCI

NCI is governed by ANU on advice from the NCI Steering Committee which serves as its Board, with ANU providing the legal framework to give effect to the independence provided for in the agreement under which NCI was established.

Its composition, which includes representation from the leading national research organisations which have a substantial interest in the advancement of the nation's high-end computing services, is as follows:

- Independent Chair—Emeritus Professor Mark Wainwright, AM
- Director, NCI—Professor Lindsay Botten
- Representatives of:
  - The Australian National University—Professor Robin Stanton, Pro Vice-Chancellor (e-Strategies)
  - CSIRO—Dr Alex Zelinsky, Group Executive, Information Sciences
  - Bureau of Meteorology—Dr Neville Smith, Deputy Director (Research and Systems)
  - Geoscience Australia—Dr Chris Pigram, Deputy CEO
  - VLSCI, University of Melbourne—(to be advised), Director, VLSCI
  - Research Intensive Universities—Prof Les Field, Deputy Vice-Chancellor (Research), UNSW (and Chair, Go8 Committee of Deputy Vice-Chancellors (Research))
- Chair, NCI Merit Allocation Committee—Professor Brian Yates, University of Tasmania

Through ANU, NCI signals its intention to rename the Steering Committee as the NCI Board from 2010 onwards.

### 4.2 Management of NCI

The management structure of NCI has not changed since its establishment in 2007. Since then, the reach of NCI, and the span of its activities has grown considerably, as has the Director's workload.

The NCI Steering Committee has agreed in principle to the appointment of an Associate Director and the Director has been requested to develop a revised organisational structure and a duty statement for the Associate Director.

### 4.3 EIF Project Governance

At its 9 March 2010 meeting, the NCI Steering Committee put in place the formal planning and project governance framework for the development of the EIF Climate HPC System.

ANU will govern the project on advice from the NCI Steering Committee, reporting to the Commonwealth Government through the designated ANU Contract Delegate (Professor Robin Stanton).

#### 4.3.1 Project Steering Committee

The Project will be overseen and directed by a Project Steering Committee established by the NCI Steering Committee.

The Project Steering Committee will comprise executive representation drawn from the other major stakeholders—CSIRO and the Bureau of Meteorology (BoM)—of this project, and will be chaired by the ANU Contract Delegate, a member of the Executive of the Australian National University. The Director, NCI will be Deputy Chair of the Project Steering Committee.



Carriage of particular aspects of the Project will be undertaken by three Planning Groups, the chairs of which will also be members of the Project Steering Committee.

#### *Membership*

Prof. Robin Stanton, Pro Vice-Chancellor (e-Strategies), ANU (Chair)  
Prof. Lindsay Botten, Director, NCI (Deputy Chair)  
Dr Alex Zelinsky, Group Executive, Information Sciences, CSIRO  
Dr Neville Smith, Deputy Director (Research and Systems), BoM  
Mr Allan Williams, Director, IT Infrastructure, Division of Information, ANU  
Dr Ben Evans, Manager, National Facility, NCI

### **4.3.2 Planning Groups**

The roles and membership of the three planning groups are summarised below.

The first two of the planning groups will commence their formal work immediately on signing of the contract, while the work of the third will commence in the latter half of 2010.

#### **4.3.2.1 Services and Infrastructure Planning Group**

##### *Role*

To advise the Project Steering Committee on:

- the scope of the science challenges and national benefit requirements that can be reasonably met within the available funding envelope for infrastructure,
- the services (i.e., computational modelling, data storage, data services and visualisation, and user tools) that are required by the priority area (both research and user communities), and by other key research communities that are major users of NCI facilities and services,
- developments in current and future high-end computing systems and data storage, through engagement with vendors and relationships with peer international organisations,
- a high-level specification of the system infrastructure.

##### *Membership*

Prof Lindsay Botten, NCI (Chair)  
Prof Andy Pitman, UNSW  
Prof. Christian Jakob, Monash University  
Professor Nathan Bindoff, University of Tasmania  
Dr Kamal Puri, Bureau of Meteorology  
Dr Tony Hirst, CSIRO  
Dr Alf Uhlherr, CSIRO  
Dr Ben Evans, ANU / NCI  
Mr Tim Pugh, Bureau of Meteorology  
Mr Chris Johnston, Department of Climate Change  
Dr Adrian Burton, ANU / ANDS  
Prof Jean Palutikof, Griffith University



#### 4.3.2.2 Data Centre Planning and Implementation

There is a natural separation between the specification of the requirements of the data centre and its subsequent detailed design and construction, imposed both by the nature of the tasks and also by the management of construction projects under the Major Project Governance Framework at ANU.

Accordingly, there will be a:

- Data Centre Planning and Implementation Group,
- Data Centre Construction Project.

##### Data Centre Planning and Implementation Group

###### *Role*

To develop a high-level specification of the data centre, taking into account the ICT requirements, anticipated future needs, contemporary best practice internationally, thereby informing the overall design of the building's fabric and services.

###### *Membership*

Mr Allan Williams, Director, IT Infrastructure, Division of Information, ANU (Chair)  
Mr Brian Lane, Associate Director, Facilities and Services, ANU  
Mr David Toll, Chief Information Officer, CSIRO  
Prof Lindsay Botten, Director, NCI  
Dr Ben Evans, Manager, National Facility, NCI

##### Data Centre Construction Project

The governance framework for major construction projects at ANU requires a:

- Project Steering Committee, which has high level oversight of the project,
- Project Management Group, which serves as the reviewing body for business decisions and processes related to the project.

###### *Membership (Construction Project Steering Committee)*

Prof Robin Stanton, Pro Vice-Chancellor, ANU (Chair)  
Mr Mick Serena, Director, Facilities and Services, ANU  
Prof Lindsay Botten, Director, NCI  
Mr Allan Williams, Director, IT Infrastructure, Division of Information, ANU  
Project Coordinator, Construction Company (to be appointed)

###### *Membership (Construction Project Management Committee)*

Mr Brian Lane, Associate Director, Facilities and Services, ANU (Chair)  
Mr Allan Williams, Director, IT Infrastructure, Division of Information, ANU  
Dr Ben Evans, Manager, National Facility, NCI  
Mr Geoff Barlow, Data Centre Facilities Manager, Division of Information, ANU  
Project Coordinator, Construction Company (to be appointed)



### 4.3.2.3 System Planning and Implementation Group

#### *Role*

The System Planning and Implementation Group will:

- develop a specification of a balanced petascale system that will meet the particular requirements of the climate change, the earth system sciences and national water management priority area, and simultaneously accommodate the breadth of high-level needs of other research communities,
- oversee the development of comprehensive tender document,
- oversee the development of a detailed evaluation plan,
- oversee the Request for Proposals process,
- serve as the Evaluation Committee, overseeing the work of the technical, financial and corporate, and legal evaluation teams, compiling a recommendation for consideration by the Project Steering Committee, and a subsequent recommendation to ANU by the NCI Steering Committee.

#### *Membership*

Dr Ben Evans, Manager, National Facility, NCI (Chair)

Prof Lindsay Botten, Director, NCI

Mr Allan Williams, Director, IT Infrastructure, Division of Information, ANU

Dr David Singleton, Manager (HPC Systems), National Facility, NCI

Dr Alf Uhlherr, Executive Manager, e-Science, CSIRO

Mr Tim Pugh, Senior ICT Officer, BoM

Prof Brian Yates, Chair, Merit Allocation Committee, NCI



## 5 Activities to Progress the NCRIS and EIF Projects

This section outlines activities to advance the NCI program in both the NCRIS and EIF funded activities. Each heading references the source of funds (in a parenthesised entry) by which the initiatives will be pursued.

### 5.1 Peak Systems — Planning, Acquisition and Commissioning (NCRIS/EIF)

#### 5.1.1 General Purpose GPU Trial (NCI CT&T, NCRIS)

With increasing interest in, and uptake of, processor accelerator technology, and particularly the general purpose graphical processor units (GPU), NCI will implement a pilot project to explore the application of accelerator technologies during 2010.

It is proposed to install a small trial system in 2010Q2 to explore targeted applications and scope possible investment in the building out of *vayu*, the current peak system in the latter half of 2010.

##### *Key Tasks*

- 2010Q2-Q3: GPU pilot project, scoping applications and future investment in accelerator technology; acquisition of small trial system (approx. \$100K); support resources provided from current NCI CT&T program.

#### 5.1.2 Expansion of the Current Peak System (EIF)

The current peak system, *vayu* (a Sun/Oracle Constellation), will be commissioned in full by mid-April 2010, with three-quarters (3/4) of the ultimate capability having been commissioned and made available to the research community on 13 March 2010.

With EIF funds, NCI proposes to enhance the capability of the system, particularly to increase the resources available to the priority access area (i.e., climate change, earth systems science) and to flagship projects which are evolving through discussion with leading research communities and research centres.

The performance enhancements that are planned are intended to increase peak performance (using conventional CPU technology) by around 30% to approximately 200 Tflops.

Dependent upon the outcome of the accelerator technology pilot study, NCI may decide to procure GPU technology to augment the peak system, with this being dependent upon the demonstrated utility of the technology for NCI's suite of research applications, the performance advantages obtained, the ease by which these may be realised by users, and the overall demand for this service.

##### *Key Tasks*

- 2010Q1/2: Initiation of necessary procurements under the current contract with Sun Microsystems.
- 2010Q3: Decision point to incorporate significant GPU technology in *vayu* based on the outcomes of the pilot project.
- 2010Q3: Installation of additional CPU capacity on *vayu*—increasing performance to 200 TFlops (approx).
- 2010Q4: Possible acquisition of GPU capability to augment the performance of *vayu*.

#### 5.1.3 Storage Infrastructure (EIF)

During calendar year 2010, NCI plans to commission storage — disk and tape infrastructure — to provide for the initial needs of the planned Australian node of the international Earth System Grid (ESG). Ultimately, the ESG node will require 2-3 Pbytes towards the end of the period 2012-13, in the lead up to the 5<sup>th</sup> assessment round of the IPCC (AR5).



During 2010-11, sufficient disk storage and new tape robots to meet the needs that exist in 2010-11 will be acquired from EIF Climate HPC funds. In subsequent years, NCI will seek to augment this initial investment in storage for what is a nationally and internationally significant dataset through alternative sources of funds, including that which will become available through a parallel Super Science initiative funded to provide for such infrastructure.

#### *Key Tasks*

- 2010Q2: Tender for data storage infrastructure (disk, tape robot).
- 2010Q3: Installation of data storage infrastructure.

### **5.1.4 EIF Climate HPC System Infrastructure and Services Planning**

#### **5.1.4.1 Services and Infrastructure Planning (EIF)**

As described in the Project Plan, the procurement of the petascale system infrastructure is underpinned by a comprehensive planning exercise, involving leading researchers and practitioners from Australia's climate change, earth systems science and water management communities, to define the required services and infrastructure using a top-down analysis of the research and national benefit drivers. This, together with consultation with leaders of other key research communities, and a parallel engagement with vendors of HPC systems, will lead to a design specification that will meet the broad requirements of stakeholder communities.

More broadly, the key stakeholders in the earth system sciences (BoM, CSIRO, University Climate Consortium, Department of Climate Change) have identified the need for a (strategic) decadal plan to evolve and meet the research and modelling services required by this community in coming years. To support the development of this strategic framework, it is envisaged that NCI will provide tactical and operational planning input to this task through the work of its Services and Infrastructure Planning Group (SIPG).

The specification of services by the SIPG involves input from the climate change, earth systems science and water management communities, and will include input from other leading computational scientists and also from ANDS (Australian National Data Service). Ultimately this will inform the infrastructure design specification (compute, data storage, and visualisation, expertise) to be procured under this contract and, in turn, the development of a system tender and also a refined estimate of the recurrent costs.

This planning task commenced informally in 2009Q3 and will be completed by September/October 2010 under the (formalised) SIPG established by the NCI Steering Committee. During the course of this work, NCI will be working with proponents of a "shortlisted" application for an ARC Centre of Excellence (CoE) application in climate science that is being prepared by partners of the University Climate Consortium, and also the proponents of other shortlisted ARC CoE applications for which high-end computing is a key requirement.

#### *Key Tasks*

- 2010Q3/Q4: Completion of the services and infrastructure plan by the Services and Infrastructure Planning Group.

Membership of the Services and Planning Infrastructure Group which oversees planning in this area is listed in Sec. 4.3.2.1.

#### **5.1.4.2 Petascale System Acquisition and Commissioning (EIF)**

This phase of the project builds directly on the work of the Services and Infrastructure Planning Group (Sec. 5.1.4.1) and undertakes the work required to establish a detailed system specification, and to prepare a comprehensive Request for Proposal and concomitant Evaluation Plan.

These documents, the drafting of which will benefit from the experience gained through the recent joint ANU-BoM procurement, will be developed by staff from the NCI National Facility, CSIRO and the Bureau of Meteorology, with reference input being provided by leading computational scientists.



This phase will also include contract negotiations with the successful vendor—with the ANU contract negotiating team being led by ANU EIF Contract Delegate (a member of the ANU Executive), and including the Director, NCI and the Manager, NCI National Facility.

#### *Key Tasks*

- Oct. 2010–Jan. 2011: Development of a detailed Request for Proposal and Evaluation Plan (NCI/ANU, CSIRO and BoM staff).
- Feb. 2011 – June 2011: Public procurement process (Request for Proposal) and evaluation of proposals.
- July 2011 – September 2011: Contract negotiations and placement of an order.
- Jan. 2012 – June 2012: Petascale HPC system installation and commissioning.

Membership of the Systems Planning and Implementation Group which oversees planning in this area is listed in Sec. 4.3.2.3.

## **5.2 Data Centre — Planning and Construction (EIF)**

A critical task for the 2010–11 financial year will be the specification, design and commencement of construction for a new data centre to:

- accommodate two high-end HPC systems (commissioned and replaced on an alternating timescale of 2+2 years), and
- be capable of serving as a major national data storage facility.

Following the announcement of the Super Science Initiative for the EIF Climate HPC System, ANU commenced the process of identifying site options for the Data Centre and engaged specialist consultants in 2009Q4 to examine options, including green field sites.

The specification of requirements of the data centre has continued in parallel with the investigation of site options and has involved comprehensive and ongoing discussions with leading HPC vendors (Sun/Oracle, SGI, Cray, and IBM) visits to state-of-the-art data centres in the USA and the UK, respectively in 2009Q4 and 2010Q1.

At the time of the preparation of this Business Plan in March 2010:

- a high level specification of the peak system infrastructure, sufficient to characterise the capability and performance, system specifications, power consumption and space requirements, within the funding envelope for a 2012 H1 installation has been developed,
- ANU has initiated a further consultancy to examine ANU campus site options and to develop a high level design specification for the data centre (consistent with the system specifications referred to above) that is to accommodate two generations of peak computing and data infrastructure, be capable of being upgraded and/or expanded for up to fifteen (15) years, and which can be constructed and commissioned by early 2012.

It is anticipated that ANU will advise DIISR by mid-May 2010 of the recommended site for the Data Centre construction.

More complete information about the carriage of the construction project will be provided to DIISR as the design and planning phase evolve during 2010–11.

Membership of the Data Centre Planning and Implementation Group which oversees planning in this area is provided in Sec. 4.3.2.2.

#### *Key Tasks*

- 2010Q2: Decision by NCI on the site of the Data Centre—to be formally notified to DIISR, triggering a project payment.
- 2010Q1-Q4: Detailed specification, design, planning and approvals for the Data Centre.
- 2010Q4: Selection of construction contractor.



- 2010Q4: Notification to DIISR of the planned commencement starting date for building construction—triggering a payment.
- 2010Q4: Initiation of data centre construction.

### 5.3 Specialised Facilities (NCRIS)

#### 5.3.1 Background

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.

To meet the need for such requirements, NCI has allocated \$2.4M to be used for acquiring shares in computational / data services available through other specialist providers, with the access and services acquired being made available through the NCI Merit Allocation Scheme.

As reported in the previous Business Plan (2009–10) and Annual Report (2008–09), the NCI Steering Committee recommended investment in two such facilities:

- Specialised Facility in Bioinformatics based at the University of Queensland,
- Specialised Facility in Imaging and Visualisation based at Monash University.

#### 5.3.2 Present Status

Contracts for the two specialised facilities were negotiated in the latter half of the 2009 and were formally executed with their host organisations in the week commencing 21 December 2009.

The University of Queensland and partners has recently finalised a contract with SGI for the procurement of infrastructure, while Monash University and partners will shortly release a Request for Proposals.

In both cases, services are scheduled to commence in mid-2010. However, delays in the procurement and commissioning of infrastructure may delay the commencement of services until August 2010.

At time of preparing this Plan, NCI is working with each of the Principals to set in place access to the services through the NCI Merit Allocation Scheme.

The following tables summarise, for each of the Specialised Facilities, the schedule of investments, the provision of services, the co-investing partners and the level of coinvestment, and the services to be provided.

#### 5.3.3 Services: Specialised Facility in Bioinformatics

<i>Contract Signatories</i>	University of Queensland, QCIF
<i>NCI Investment</i>	\$1.2M over two years (2010H1-2011H2)
<i>Total Investment (incl. NCI)</i>	\$5.657M (over 3 years) \$3.150M cash, \$2.507M in-kind
<i>NCI Resource Share</i>	Ave: 21.2%
<i>Period of Service Provision</i>	Three years: 2010H2-2013H1 (services valued at \$300K p.a.)
<i>Partner Organisations</i>	UQ, QCIF, CSIRO, NCI
<i>Services to be provided</i>	<ul style="list-style-type: none"> <li>▪ HPC Infrastructure, Data Storage, Support</li> <li>▪ Comprehensive range of databases and software including: GenBank, RefSeq, UniProt, OMIM, Reactome, UniGene, Enzyme, KEGG, PDB, interPro, GO, UniGene;</li> </ul>



	<ul style="list-style-type: none"> <li>▪ Access to dynamic mirrors of key international data in genomics and bioinformatics including the proposed Australian EBI Mirror, and the two main genome browsers (Ensembl, USCS) via SRB, integrated where possible via an integration tool such as SRS web services;</li> <li>▪ Access to a range of public-domain software tools including:             <ul style="list-style-type: none"> <li>○ Plone adapted for use with SRB stores</li> <li>○ SRB/iRODS and Hermes (for alignment with ARCS and ANDS)</li> <li>○ the myGRID/Taverna workflow environment</li> <li>○ NIMROD and associated tools for grid computing and parameter estimation QosCosGRID middleware;</li> </ul> </li> <li>▪ Access to commercial software packages including:             <ul style="list-style-type: none"> <li>○ SRS</li> <li>○ Inforsense Biosense</li> <li>○ GeneGo</li> <li>○ Spotfire DecisionSite.</li> </ul> </li> </ul>
--	--

### 5.3.4 Services: Specialised Facility in Imaging and Visualisation

<i>Contract Signatories</i>	Monash University
<i>NCI Investment</i>	\$1.2M over two years (2010H1-2011H2)
<i>Total Investment (incl. NCI)</i>	\$8.376M (over 3 years) \$5.550M cash, \$2.826M in-kind
<i>NCI Resource Share</i>	Ave: 14.3%
<i>Period of Service Provision</i>	Three years: 2010H2-2013H1 NCI services valued at \$300K p.a.
<i>Partner Organisations</i>	Monash, Australian Synchrotron, VPAC, CSIRO, NCI
<i>Services to be provided</i>	<ul style="list-style-type: none"> <li>▪ HPC Infrastructure, Data Storage, Support</li> <li>▪ Comprehensive range of imaging and visualisation software:             <ul style="list-style-type: none"> <li>○ Volume reconstruction, segmentation and visualisation: X-TRACT, Insight Tool (ITK), Imaris, Seg3D, VolView, Drishti;</li> <li>○ Imaging processing and analysis: Matlab, Matlab Image Processing and Wavelet Toolboxes, IDL, CImg, ImageJ Open Computer Vision Library, BioImage Suite;</li> <li>○ Visualisation tools and libraries: Visualisation Toolkit (VTK), OpenDX, Ensign Gold, Visit, Paraview, Vis5D+, SciRun, ImageVis3D, MayaVi, Chromium or Powervue;</li> <li>○ Domain specific visualisation and image processing libraries: FSL and SPM, MedINRIA, DtiStudio.</li> </ul> </li> </ul>

## 5.4 Specialised Support and Development

Support and development activities that are tailored to the needs of research communities are important in advancing the uptake and sophistication of usage of Australian high-end computing, and also in increasing the impact of research for which high-end computing is a critical element.



The Specialised Support and Development Program is to be the successor to the Computational Tools and Techniques Program (under the NCRIS NCI project), operating under a more appropriate name.

It includes three existing CT&T activities:

- Specialised computational support for the earth systems science research community,
- Specialised computational support for the astronomy community,
- Implementation of support for data-intensive computation based on the cloud computing paradigm,

and introduces one new activity, a

- Development program, to enhance the scaling and efficiency of high-impact codes implemented on the National Facility.

The current status of each activity, and plans for their development in the period covered by the Business Plan, are outlined below (which also note the funding stream for each activity).

#### 5.4.1 Earth Systems Science (NCRIS CT&T)

The ESS support program has been developed in consultation with leaders of the ESS research community to provide for closer integration of the university climate research community with their counterparts in CAWCR (CSIRO and the Bureau of Meteorology), in order to maximise the benefits of the petascale system when it is commissioned in 2012.

There are two full-time support positions provided for in each of the calendar years 2010 and 2011 to meet the objectives of:

- enhancing the liaison between the Australian University climate community, CAWCR and the Hadley Centre (UK Met Office) in the development and uptake of the ACCESS model;
- providing for the building, maintenance and continuous upgrading of the ACCESS model (version by version) plus support for a broader range of climate community needs on the NCI National Facility;
- designing, implementing and maintaining shared data repositories (and associated workflows) directed towards the development of the Australian node of the international Earth Systems Grid, together with the efficient analysis of that data.

One position is based at the NCI National Facility, while the other is located in Melbourne, the hub of the development of the ACCESS code, and is hosted by Monash University. A staffing agreement to provide for the Monash based position was signed with Monash University in late 2009, with Dr Michael Rezny, formerly with SGI, having taken up the position recently. The NCI National Facility position will be filled shortly.

Milestones for the ESS support program, which are encapsulated in the project plan, are listed below. A project review is planned for late in 2010.

##### 2010Q1

- Implementation of ACCESS system on the new NCI National Facility,
- First basic user support system in place.

##### 2010Q2

- First draft of a strategic plan for future development of the ACCESS system at NCI developed in collaboration with CAWCR/Hadley Centre.

##### 2010Q3

- Implementation of joint subversion server,
- First user workshop,
- Stakeholder workshop (with ACCESS).

2010Q4

- Final draft of the strategic plan(with CAWCR/Hadley Centre),
- Unified file conversion software implemented and ready for use.

2011Q1

- Project review —2011Q1.

2011Q2

- Second user workshop,
- Proposal for possible extension of the project developed and submitted.

2011Q3

- Update of the strategic plan for further development,
- Stakeholder workshop — Q3 2011.

General

- Implementation of ACCESS system updates — ongoing as needed.

### 5.4.2 Astronomy (NCRIS CT&T)

The Astronomy support program has been developed in consultation with leaders of the national astronomy research community through Astronomy Australia (and particularly its HPC Working Group).

Again, there will be two full-time positions supported for two years, with the agreed plan providing for one at the NCI National Facility and the other at the Swinburne University Centre for Astrophysics and Supercomputing (CAS).

The main objectives of the support program are to:

- Enhance the effective utilisation of national supercomputing facilities in astronomy and astrophysics,
- Advance and facilitate the analysis of large astronomical data sets.

The NCI National Facility and Swinburne based positions will respectively have foci in optical and radio astronomy, with both advancing the modelling capability available to Australian astronomers and astrophysicists. The modelling will be undertaken on the NCI National Facility and the computational facility at the Swinburne CAS, with the latter being augmented with substantial general purpose GPU capability to be installed with funds from Swinburne University and Astronomy Australia (through a Super Science Initiative).

The projects that are planned are:

Swinburne CAS:

- Porting and optimisation of codes for CPU-only and CPU/GPU use on the Swinburne CAS systems,
- Supporting the high time resolution survey of the universe (pulsars).

NCI NF

- Porting and optimisation of astronomy codes which scale to high core counts,
- Supporting data analysis and visualisation of large observational and simulated data sets.

Milestones for each project are listed in the Project Plan.

The staffing agreement to provide for the Swinburne position should be executed in early April 2010, with both appointments being in place in 2010 Q2.



### 5.4.3 Cloud Computing Services (NCRIS CT&T)

#### Background

The typical supercomputer provides a high performance computational environment for work that requires raw computational power and low latency access to memory and other computational nodes. In contrast, data intensive computation typically focuses on high throughput operations to be applied to terabytes (or even petabytes) of data, but without the expectation of low latency, highly connected bandwidth. This environment has evolved to meet the search, processing and analysis needs imposed by massive datasets generated by powerful new instruments, networks of remote sensing equipment, and accessible large-scale data repositories.

This new style of computing presents a range of interesting new imperatives and challenges including programming models, parallel algorithms, and applications, along with the underlying service delivery, system design and resource management. NCI, in preparing for future data-intensive services, on both local infrastructure and in the commercial computing clouds, has a development program underway to explore the new paradigm, comparing the facilities and services available in commercial clouds (particularly the Amazon Elastic Cloud EC2) with freely available implementations (Eucalyptus Elastic Cloud and VMware Cloud APIs) that are installed on the National Facility Data Cluster ([dc.nci.org.au](http://dc.nci.org.au)).

#### Application Areas

NCI's initial efforts involve three application areas: Astronomy, Climate and Environmental Science, and Bioinformatics.

##### *Astronomy*

The astronomy data generated from optical telescopes such as SkyMapper require next generation search techniques with capabilities beyond those provided for in relational databases. The Map Reduce programming model is being trialled for projects such as the LSST (the Large Synoptic Survey Telescope) for possible future implementation through a cloud service. In order to facilitate this work for SkyMapper, it is necessary to enable the Hadoop on Demand service for the peak system ([vayu.nci.org.au](http://vayu.nci.org.au)) and the data cluster ([dc.nci.org.au](http://dc.nci.org.au)).

##### *Climate and Environmental Science*

The CMIP5 data being generated by computational models in preparation for the IPCC assessments are estimated to reach approximately 12 PBytes in size by the end of 2013. This dataset will be used by climate scientists throughout the world and will be made available to Australian researchers through the EIF Climate HPC Project and the parallel EIF Research Data Storage Initiative. Such data must be highly available and be able to be explored and analysed using both traditional and new data analysis techniques for which the data cloud is ideal.

##### *Bioinformatics*

Crossbow is a scalable pipeline for whole genome resequencing analysis, combining Bowtie, an ultrafast and memory efficient short read aligner, and SoapSNP, an accurate genotyper, within Hadoop to distribute and greatly accelerate the computation across many nodes in a cloud environment. The NCI cloud team is implementing Crossbow on NCI systems and, as a prototype service, it is attempting to re-run the SNP analysis reported in the Crossbow paper (<http://genomebiology.com/2009/10/11/R134>). To facilitate this, the Yanhung genome dataset has been downloaded from the BGI.

#### Key Activities for 2010–11

The key tasks and milestones for the coming year are:

- Deployment of HDFS filesystem and its population with IPCC data to be served using ESG (Earth Systems Grid) data node software. This includes integration with fast data transfer protocols such as gridFTP,
- Deployment of Hadoop on Demand infrastructure,
- Participation in the SC'10 bandwidth challenge with the international climate supernodes involved in the ESG,



- Comparison of the functionality and cost effectiveness of commercial Elastic Clouds (e.g., Amazon EC2) with similar environments at the National Facility.

#### 5.4.4 Code Scaling Project (ANU/NCI-Oracle-Collaborative Agreement)

At the cutting edge of research, high performance computing has long been a key driver of scientific discovery. For many years, until comparatively recently however, the available computational capabilities have lagged behind the research agendas. However, with the evolution of petascale architectures, and ultimately the next generation of exascale systems, this situation is becoming transposed in a number of research communities. This is certainly the situation in Australia—a matter that is under the constant attention of the NCI Steering Committee which understands that the transformative potential of high-end computational infrastructure on cutting edge research will not be realised without substantial efforts that increase the scope, uptake and sophistication of usage by Australian researchers.

In order to ensure that leading Australian computational research is prepared for the current and next generation computing environments, and is able to take full and comprehensive advantage of this infrastructure, NCI will implement a development program, the object of which is to enhance the scaling of codes on high terascale and petascale environments.

This “development program” will provide for the:

- restructuring and redevelopment of existing codes to enhance their scalability,
- investigation of better scaling codes for research communities and the porting of such codes to the National Facility,
- the development of new tools, where needed,
- exploration of emerging technologies, e.g., accelerator technologies.

The first three of these objectives will be initiated through the *Code Scaling Project* while the fourth of the objectives has been described previously in Sec. 5.1.1, and is supported (i.e., staffing) through the CT&T budget.

In the first instance, the *Code Scaling Project* will be supported by a full-time specialist programmer, appointed for three years with the position supported through the ANU/NCI-Oracle Collaborative Agreement (under the NCI National Facility contract between ANU and Oracle (previously Sun Microsystems)). Under the ANU/NCI-Oracle agreement, there will also be access to specialist expertise with Oracle that will augment and add value to this local capability.

Beyond that, the imminence of the installation of the petascale system in 2012 necessitates a significant increase in the size of the support base at NCI. With current workloads, the existing support team is largely committed to servicing the Help Desk, with little opportunity to engage in developmental activities that would increase the effectiveness or quality of usage for users. In preparation for the petascale resources that will become available in 2012, there is an urgent need to increase the scope and the impact of the support provided by NCI.

This developmental focus, which is further amplified in Sec. 6.1.2, will be crystallised through the creation of a Development Team within the National Facility, the focus of which will be on advanced code scaling, code development, data-intensive computation, advanced data management etc. The Budget provides for four positions in 2011, rising to eight positions in 2012, funded initially through the NCRIS NCI project and the ANU/NCI Collaborative Agreement. The decision to proceed with this initiative, at the proposed level of commitment, is consequent upon the provisions of the Collaborative Agreement which, beyond the residual NCRIS funding, will need to provide for the entire recurrent cost in the EIF environment.

## 5.5 Access (NCRIS and EIF)

### 5.5.1 Background and Context

Access to NCI resources has been, and continues to be provided in two ways, through:

- Merit and Priority Access,



- Partner Allocations, in which NCI partners and affiliates acquire resources (respectively a defined share or a fixed quantity of resources).

#### 5.5.1.1 Merit and Priority Access

From the inception of NCI, the former, i.e., Merit and Priority Access, has been provided for through NCRIS funding which provides the Commonwealth share of the resources on the National Facility, and, from 2010, of resources acquired in each of the two Specialised Facilities.

During the period of the current Business Plan (2010–11), however, there will be changes to the operation of the Merit and Priority Access Scheme, specifically to provide for the:

- priority access field (climate change/earth systems science and national water management) associated with the EIF objectives, and the
- flagship allocation scheme, detailed in the previous Annual Report.

#### 5.5.1.2 Partnership (Research Organisations)

Since the inception of NCI in 2007, there has been an expectation of substantial co-investment by research organisations and consortia, complementing and supplementing the Commonwealth investment, in order to develop both the facilities and the services provided to researchers.

This expectation has been realised and is encoded in the NCI Partner Service Agreement which provides for Partner and Affiliate Membership (both of which are described in previous Business Plans and Annual Reports). Indeed, in 2010, the annualised partner investment in NCI will reach approximately \$7.5M per annum.

At present, i.e., under the current NCRIS NCI project, partner investments provide NCI with additional resources with which to develop the facilities and services in a manner which:

- benefits all stakeholders,
- provides additional certainty in NCI's financial planning (through the extended nature of partnership contracts) and, in return,
- benefits partner organisations by providing them with access to increasing levels of services and resources (that evolve through planned system increments).

In the extension of NCI under the EIF Climate HPC contract, partnership will play an even more critical role providing, as it must, for the operating costs of the facility and the development of services that are required by the priority area, and other national research priorities.

The NCI Steering Committee recognises that partner investments are made by research organisations in support of their strategic research priorities and requirements. Accordingly, the Steering Committee acknowledges that the resources and services made available to partner organisations under the NCI Resource and Services Allocation Model will be theirs to prioritise in support of their strategic research mission.

#### 5.5.1.3 Industry Partnership and Commercial Access

Consistent with the current innovation initiatives of the Commonwealth Government, the NCI Steering Committee has agreed to an extension of the concept of “partnership” to accommodate industry and commercial relationships.

The industry/commercial partnerships will:

- assist in fostering the uptake of high-end computing in industry,
- provide a mechanism by which NCI can build a portfolio of industry engagement, and
- generate additional income with which to advance NCI's suite of services.

This will be formalised in a comprehensive Commercial Access policy to be developed during 2010.



## 5.5.2 Access Schemes

### 5.5.2.1 Priority Access

During 2010Q2-2010Q3, the Services and Infrastructure Planning Group (Sec. 4.3.2.1) will identify the broad requirements of the climate change, earth systems science and national water management community, while a similar exercise will identify the corresponding needs in other research areas.

In some cases, partner shares will be used in support of the climate change, earth systems science and the national water management priority area, while in others it will be used to support other institutional research strengths and institutionally prioritised projects, or some combination thereof.

The Commonwealth share will be allocated:

- in support of the priority area,
- to other targeted, high-impact research fields, through the flagship allocation scheme, and
- to support the breadth of high quality research (already sustained by national research granting agencies) that requires access to high-end computing services.

Principles underpinning the allocation of resources in the priority area will be developed by the Project Steering Committee and will be endorsed subsequently by the NCI Steering Committee before being reported to the Commonwealth as a project deliverable through annual business plans and annual reports.

It is intended that these revised allocation principles will underpin the 2011 round of resource allocation to take place late in 2010 for the newly augmented (with EIF funds) NCI National Facility peak system.

### 5.5.2.2 Flagship Access

The previous Annual Report (2008–09) foreshadowed the introduction of the Flagship Allocation Scheme, the key goals of which are to:

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

NCI is aiming to support computationally intensive research of the highest impact, by providing flagship projects with a package of resources and support that provides for:

- internationally competitive allocations,
- certainty of allocation through multi-year grants (two years in the first instance), and
- high levels of developmental support, where appropriate.

This additional support is intended to add value to the research and to raise its impact and ambition through a range of activities including enhancements of the scaling of existing codes, porting new (better scaling) codes to the peak system, assisting with the management and visualisation of large data sets, and so forth.

The flagship allocations will build on the support programs already in place including the:

- existing National Facility Support Team,
- specialised support that is provided for the earth systems sciences—the priority allocation area— (Sec. 5.4.1), and astronomy (Sec. 5.4.2) communities, and particularly,
- new Code Scaling Program (Sec. 5.4.4), referred to previously,



- envisaged Development Team (if the initial budgeted allocation can be sustained through co-investment under the Collaboration Agreement).

At the time of preparation of this Business Plan, NCI is working with a number of shortlisted applicants for the current (2010) ARC Centres of Excellence round, having written to their proponents (though a letter circulated by the ARC Executive Director with carriage of the ARC Centres program) expanding on NCI services and offering to assist and support their applications. It is likely that NCI will be supporting at least three such applications through its role as a formal “partner organisation”.

Centre of Excellence (CoE) applicants understand that the promised entitlement under the Flagship Scheme, initially for a period of two years, will be honoured if the application is successful and that, by virtue of their success, no further application to NCI will be required—the NCI merit criteria being deemed as met. Applicants further understand that they need to adhere to NCI’s annual reporting requirements. In the event that a supported CoE application is not successful, applicants understand that NCI may review any commitment that was made.

Expressions of interest may be sought during 2010 to further develop the Flagship Scheme, consistent with the criteria reported in the previous (2008–09) Annual Report from NCI, viz.:

- Demonstrated science excellence and research standing,
- A persuasive case for support under the Flagship Program, e.g., superior research questions / outcomes, higher impact results, deeper research questions posed / answered, enhanced sophistication of the use of computational processes leading to superior outcomes etc.

### 5.5.2.3 Open Access Merit Scheme

The Open Access Merit Allocation Scheme is the continuation of the current NCI Merit Allocation Scheme and will provide for the diversity of high-quality research that requires access to computational modelling on high-end systems.

Its operations will continue under the same merit criteria as presently exist, and which have been provided in previous Business Plans and Annual Reports.

### 5.5.3 Allocations for Calendar Year 2010

Scheme/Partner	Share (%)	Share (M core hrs)
MAS	44.093	46.135
ANU	25.394	26.570
CSIRO	25.020	26.179
Intersect	4.186	4.380
Monash	0.560	0.586
iVEC	0.373	0.391
QCIF	0.373	0.391
Total	100.000	104.631

At its November 2009 Meeting, the NCI Steering Committee agreed that approximately 95% of the resources available through the Sun Constellation (vayu) and the SGI Altix XE (xe) would be allocated to the Merit Allocation Scheme and Partner Shares, with approximately 5% being retained as the Director’s Share to accommodate unforeseen requirements and new projects, including possible commercial access.

Approximately 104.6M core hours were allocated under the current Resource Allocation Model according to the distribution in the table on the previous page.



For the Merit Allocation Scheme, the Steering Committee further decided that:

- 70% (i.e., up to 32.27M core hours) of the notional allocation to the MAS should be available for allocation under the Open Access Scheme by the Merit Allocation Committee at its December 2009 meeting,
- 30% should be retained to accommodate the new Flagship Allocation.

#### Brief Report of 2010 Allocations for the Open Access MAS Scheme

- A total of 25.1 MSU was allocated for 2010 by the Merit Allocation Committee (MAC), representing an under-allocation of that made available by the NCI Steering Committee. Some applications were not granted any resources under the MAS, while others were provided with less than their request, with feedback being provided from the MAC about concerns with their application that should be addressed.
- The MAC emphasised the need to ensure uniformly high quality amongst those projects that were allocated the largest share of the resources. Such projects were able to demonstrate productive use of peak computing resources, produce excellent research and publications in A and A\* classified journals, and were able to demonstrate significant impact (citation data).

- Recent year comparisons:

Year	Allocated	Requested	Available
2007	12.3 MSU	19.2 MSU	10.2 MSU
2008	13.9 MSU	19.5 MSU	12.8 MSU
2009	15.3 MSU	28.4 MSU	12.8 MSU
2010	25.1 MSU	41.0 MSU	32.3 MSU

#### **5.5.4 National Merit Allocation Scheme**

The previous Annual Report (2008–09) drew attention to the need for a nationally coordinated access scheme for the multiple high-end computing systems that are now under development in Australia. This matter was further developed in the letter from DIISR (from Mr David Luchetti dated 7 September 2009) to ANU, inviting ANU to negotiate an EIF contract. In that letter, amongst other matters, the Commonwealth sought “an indication of the allocation process that could be used to ensure world class science outputs, and the manner in which allocation processes across the various peak facilities could be harmonised.”

In keeping with best practice internationally, NCI, through ANU, advocated a nationally coordinated access framework that was referred to in the reply as a National Merit Allocation Scheme. Indeed, the Teragrid of the USA has returned to a single national allocation process and interface for access to its major facilities, having abandoned its earlier uncoordinated approach.

Since last year, when this matter was first raised in the NCI Annual Report and in correspondence with DIISR, the Victorian Life Sciences Computation Initiative (VLSCI) has established its merit allocation processes, with these adopting (with only minor tailoring) the NCI Merit Allocation Criteria. It is notable that the NCI Director has been appointed as Chair of the VLSCI Resource Allocation Committee in order to provide developmental input and to ensure consistency of process and principles.

The ANU / NCI response to the letter from Mr Luchetti outlined principles and a process whereby the NCI Merit Allocation Scheme might be extended nationally, with a single interface for applications (under open access) to the Commonwealth funded facilities (NCI and Pawsey Centre).

Key elements of this, recorded in the paper attached to ANU’s formal reply, were:

- the composition of a National Merit Allocation Committee (NMAC) being agreed by the Directors of the peak facilities (on advice from research communities), taking into account the requisite matrix structure (which combines expertise and geographical diversity of members), and approved subsequently by their Boards / Steering Committees;
- the Chair of the NMAC being nominated jointly by the peak facility Directors and approved jointly by their Boards;



- the NMAC reporting to the Boards of the peak facilities, with the MAC Chair potentially having a formal role<sup>1</sup> on each of these bodies, as is presently the case with NCI;
- the decisions of the NMAC being independent<sup>2</sup> of the management of the peak facility operators, with its decisions being implemented by each facility without variation;
- the merit assessment criteria being agreed and varied jointly through a committee comprising the NMAC Chair, and the Directors of, and a senior advisor from, each of the peak facilities, and being approved by the Boards of each of the peak facilities;
- the assessments of the NMAC being assisted by technical advice provided by staff of the peak facilities, with each of the peak facilities being invited to provide a senior technical advisor to assist the work of the meetings of the NMAC.

NCI would be pleased to work with DIISR and the Pawsey Centre to bring into effect a nationally coordinated framework to provide for access to national peak computing services.

## 5.6 Outreach Activities

The focus of NCI's outreach activities lies in promoting the profile and raising the impact of high-end computing services in Australian research. These are outlined under the relevant headings below.

### 5.6.1 Public Relations and Promotional Materials (NCRIS)

During 2009-10, NCI substantially upgraded its public face with new publicity materials, and a new website. The website upgrade continues apace and is now underpinned by a full content management system, thus facilitating editing of the site and the development of new content. During the coming year, the content of the NCI ([www.nci.org.au](http://www.nci.org.au)) and National Facility ([nf.nci.org.au](http://nf.nci.org.au)) website will be merged into the single site.

The period 2009-10 also saw the development of an attractive suite of promotional materials, the major focus of which was the "research highlights"—750 word, professionally written vignettes about high-impact, "media friendly" science. These have been universally well received as highlighting strong science outcomes and were a feature of the launch of the NCI National Facilities by the Minister, Senator the Hon. Kim Carr, on 16 November 2009.

During the coming year it proposed to:

- further develop the suite of research highlights in 2010-11, with additional materials being produced,
- exploit these through a media strategy (to be developed through a consultancy with a professional science journalist), the aim of which is to promote high-end computing and the work of NCI, and the high-impact research that it supports, in the mainstream media.

### 5.6.2 Building and Maintaining Relationships (NCRIS)

#### International Relationships

During the course of 2009-10, NCI established or consolidated a number of very useful relationships, including the EPCC and e-Science Institute (University of Edinburgh), the San Diego Supercomputer Centre (University of San Diego) and the UK Engineering and Physical Sciences Research Council. The building of further international relationships that will advance the work of NCI is planned for the period 2010–11.

#### User Survey

The quality assurance survey of NCI users planned for early in 2010 has been delayed until 2010Q3. Such a delay has been necessitated by delays in the commissioning of the Sun/Oracle Constellation. Since the full system will not be available until mid-April 2010, it is prudent to provide users with at

<sup>1</sup> For reference, the NCI MAC Chair is an independent member of the NCI Steering Committee.

<sup>2</sup> A key feature of the present NCI Merit Allocation Committee is the independence of its decision making from the management and operation of NCI. Of critical importance, is that the MAC reports to the Steering Committee, and that its allocation decisions are implemented without alteration by NCI.



least three months of services before surveying their attitudes to the services and the facilities. At this stage, however, informal surveys have indicated that researchers are very happy with the new system, the service provided, and the smoothness of the transition to it.

### **Geographical expansion of the NCI Support Team**

The expansion of the NCI National Facility Support Team from its Canberra base, anticipated in the previous Business Plan, is underway. NCI will be contributing 50 percent of the cost of a full-time staff member in each of Melbourne and Sydney, with the staffing agreement to provide for the Melbourne based position having been agreed (in late 2009) with Monash University. A further agreement with Intersect is planned in coming months.

### **Training Courses**

Workshops, training courses and seminars will continue to be provided by NCI NF staff at any eligible organisation as required. As in previous years, these activities are driven by user demand.

### **Pilot Programs**

NCI will continue to build and consolidate its community of research users and stakeholder organisations through targeted pilot programs.

The focus of activities during 2009–10 was with Geoscience Australia (GA), with this effort being rewarded by GA's recent decision (March 2010) to take up partnership with NCI.

Further pilot programs will be undertaken during 2010–11 to foster the uptake of high-end computing services, with a number of research intensive universities being targeted during the coming period.

### **User Forum**

NCI has evolved a substantial partnership over the past two years, with the organisations having signed the Partner Service Agreement by the end of the 2009–10 financial year now numbering seven (7).

In order to provide a more formal means of facilitating interaction, feedback and contributions to NCI's planning of services, NCI will establish a Partner Forum during 2010–11.

The role of the Partner Forum will be to:

- provide a formal mechanism of communication between partner organisations, their constituents, and NCI,
- advise and provide feedback to NCI on matters of service delivery,
- assist NCI in its planning of new services,
- report to the NCI Steering Committee (through the NCI Director) on the operations of the partnership,
- address other matters of relevance and significance to the management of partner relationships, as may be raised by partners and NCI.

### **5.6.3 Building Capability (ANU/NCI-Oracle Collaborative Agreement)**

High performance and high throughput computing in the petascale environment open up a range of new and exciting research possibilities, many of which necessitate, or at least would benefit from, increased multi-disciplinary collaborations.

Within the NCI community of interests that are key strengths of its major stakeholders, there are interesting possibilities in areas such as climate modelling, earth systems science, geosciences, the environment, and others, which open up new, interdisciplinary collaborations in these thematic areas.

Cognisant of the success of the "themes" program of the UK e-Science Institute in building new research capabilities, NCI will initiate and foster a program of research themes which comprise a sequence of events (e.g., workshops, visits etc.) of six to twelve months duration which bring together



leading researchers and practitioners with the aim of fostering new research collaborations and sustainable new research directions.

In the context of the planned NCI program, which is to be funded through the ANU/NCI-Oracle Collaboration Agreement (provided for within the ANU-Oracle/Sun contract), the specific aims are to:

- advance knowledge outcomes by combining discipline strengths with expertise in advanced computational and data intensive science, and so
- drive the uptake and impact of high-end computing services in these areas.

#### **5.6.4 Conferences (NCRIS)**

NCI will be represented at the annual eResearch Australasia (<http://www.eresearch.edu.au>) conference to be held in Queensland from 8–12 November 2010. This year, the Director is again a member of the conference's Program Committee.

It is important that NCI be represented at SC10 (the premier international supercomputing conference) in preparation for the release of its request for proposal for the petascale system in early 2011. Regrettably, NCI was not represented at a senior level at SC09 since the conference dates clashed with the launch of the National Facility system in November 2009.

### **5.7 Planning and Policy Development**

There are a number of significant items of planning and policy development relating to both the NCRIS and EIF activities that are required during the period covered by this Business Plan. These are outlined below.

#### **5.7.1 Industry and Commercial Access Policy**

With the introduction of the commercial/industry partnership scheme, NCI will redevelop its commercial access policy into a comprehensive policy for commercial arrangements and industry access. This policy will be developed during 2010Q2/Q3 and will be posted on the website after its approval by the NCI Steering Committee.

Also required is the development of a Service Contract to accommodate the new partnership for commercial/industry relations.

#### **5.7.2 Redevelopment of the Partner Service Agreement**

Since the inception of NCI, the growth in computation that is data-intensive has grown and so the Schedule of the Partner Service Agreement that defines the services offered by NCI needs to be made more comprehensive so as to better define and accommodate the diversity of resources and services that are now, and will be, provided.

The present Partner Service Agreement aligns with the NCRIS Funding Agreement and expires at the end of 2011. This, together with the need to put in place a Collaborative Agreement to provide for the co-investment needed to sustain the petascale system, necessitates the redevelopment of the Partner Service Agreement during the coming year.

#### **5.7.3 Resource Allocation Model**

Substantial planning efforts are needed to update the Resource Allocation Model consequent upon significant changes to the resource base.

The conclusion of the NCRIS funding agreement, and the introduction of the petascale system, underpinned by the new Collaborative Agreement, together with a significant revision of the access mechanisms to provide for priority access for climate change, earth system science and national water management, necessitates a major review (in FY2010–11) of NCI's resource costing and resource allocation model. The planned review will be undertaken during the course of the 2010–11 financial year.



The termination of the NCRIS funding line in 2011 raises concerns about the continued provision for the merit-based access schemes (i.e., the Flagship and Open Access Merit Allocation Schemes) and discussions are sought with DISSR to resolve this matter.



## 7 Milestones

All aspects of the NCRIS NCI Project (which concludes in 2011) are underway and some, including the commissioning of the new peak system, are essentially complete and now in production.

Most of the developmental activity within NCI is now associated with the EIF Climate HPC Project, with the milestones in each of the two projects (outlined under the following heading) reflecting this fact.

### 7.1 NCRIS NCI Project

#### By 30 June 2010

- Sun/Oracle Constellation commissioned in full (April 2010) — the planned extension of its capability is provided for under the EIF project.
- Completion of the staffing agreement with Swinburne University for the Specialised Support staff member in Astronomy; agreement of a project plan for astronomy / astrophysics specialised support.
- All staff appointments in place at the NCI National Facility, Monash University and Swinburne University, for the Specialised Support Programs in Earth Systems Science and Astronomy / Astrophysics.

#### By 30 September 2010

- Resources and services available through the Specialised Facility for Bioinformatics — via the NCI Merit Allocation Scheme.
- Resources and services available through the Specialised Facility for Imaging and Visualisation — via the NCI Merit Allocation Scheme.
- Completion of the development of a Commercial Access and Industry Partnership Policy.
- Initial Flagship Allocation Round (scheduled to occur after the announcement of successful ARC Centres of Excellence).
- Conduct and completion of an independent survey of the NCI user community.

#### By 31 December 2010

- Merit Allocation Round (Open Access Scheme) — conducted under the new access regime which provides for priority access under the EIF scheme.
- Revision of the NCI Partner Service Agreement, following completion of the Collaborative Agreement to underpin the support and recurrent costs associated with the EIF funded project and the ongoing expenses of the peak system commissioned under the NCRIS scheme.
- Completion of new / additional “research highlights” for the Research Gallery.
- Participation in eResearch Australasia.
- NCI attendance at SC10.

#### By 31 March 2011

- Appointment of additional code scaling expertise to the NCI National Facility Support Team (provided for through ANU/NCI-Oracle Collaborative Agreement).

#### By 30 June 2011

- Redevelopment of the NCI Resource Allocation and Budget Model.



## 7.2 EIF Climate HPC Project

The following quarterly milestones are provided for the financial year 2010–11.

Since this is the first business plan covering the EIF Climate HPC System, project milestones are also provided for the first two quarters of the 2010 calendar year.

### By 31 March 2010

- Development of the Project Plan and the initial Implementation Plan.
- Endorsement of the Project Plan and an initial Implementation Plan by the NCI Steering Committee, together with a recommendation to ANU that it negotiates and finalises a Funding Agreement with the Commonwealth (DIISR).
- Initiation and completion of a preliminary peak system infrastructure scoping exercise, sufficient to specify high level data centre requirements for 2012 and 2014-15.
- Initiation of a data centre consultancy to identify and confirm potential sites and to scope construction and related infrastructure services and costs (sufficient for ANU to confirm to DIISR a firm site location by mid-May 2010).

### By 30 June 2010

- Negotiation and execution of an EIF Funding Agreement between ANU and DIISR.
- Procurement of additional CPU capacity for the Sun Constellation system (to raise system performance to approximately 200 Tflops).
- Initiation of detailed data centre planning, design, specification, approvals, etc.
- Decision by ANU on the site of the Data Centre—to be formally notified to DIISR, triggering a project payment.
- Completion of a high-level specification of an acceptable petascale system and its associated data storage; estimation of the necessary power and cooling capacity; estimation of recurrent costs to be sustained by co-investment.
- Initiation of formal discussions between ANU, CSIRO, BoM and other co-investors leading to a Collaborative Agreement to provide for the recurrent operation of the Climate HPC Centre.
- Initiation of a tender for data storage infrastructure (disk, tape) to provision for climate change, earth systems science and national water management data, and other storage needs.
- Initiation of pilot project exploring accelerator technologies and scoping future GPU investment in NCI peak facilities.
- Initiation of an analysis of the total cost of supporting research on NCI systems (infrastructure, software and support) under an EIF contract.
- Initiation of the establishment of principles underpinning future Priority Access allocations (NCI Steering Committee).

### By 30 September 2010

- Completion of the services and infrastructure plan by the Services and Infrastructure Planning Group.
- Installation of additional data storage infrastructure (disk, tape drives) tendered for in 2010 Q2.
- Installation of additional CPU computational capacity on the Sun Constellation system.
- Completion of pilot project exploring accelerator technologies (scoping future GPU investment).



- Completion of the analysis of the total cost of supporting research on NCI systems; consideration by the NCI Steering Committee.

By 31 December 2010

- Completion of the data centre specification, design, approvals etc.
- Selection of construction contractor by ANU.
- Notification to DIISR of the planned commencement starting date for building construction—triggering a project payment.
- Notification to DIISR of completion of a Collaboration Agreement between co-investing partners, triggering an expenditure milestone with DIISR formally authorising execution of a construction contract and the commencement of construction.
- Completion of the determination of Priority Access Principles (by October 2010) in preparation for 2011 round of allocations.
- Decision and possible procurement of GPU accelerator technology to augment the NCI peak system.
- Completion of merit access and partner allocation rounds on the new (*vayu*) facilities.
- Provision of additional resource(s) in the National Facility Support Team to provide advanced support to enhance scaling of amenable research codes and advance other data intensive applications.
- Initiation of the development of comprehensive documentation for a request for proposal for the peak system infrastructure.
- Initiation of the development of the evaluation plan for the peak system infrastructure.
- Investigate and, if necessary implement, changes to infrastructure / system software to accommodate on-demand access.

By 31 March 2011

- Initiation of 12 month data centre construction and fit out (if not initiated in the previous quarter); Construction milestone updates will be provided once the construction project plan is evolved.
- Completion of the petascale infrastructure Request for Proposals and Evaluation Plan documents (February 2011).
- Initiation of the petascale infrastructure procurement (February-March 2011).

By 30 June 2011

- Continuation of the 12 month data centre construction and fit out; Construction milestone updates will be provided once the construction project plan is evolved.
- Evaluation of requests for proposals; negotiations with vendors; best and final offers; formulation of a recommendation to the Project Steering Committee and the NCI Steering Committee.



## 8 Risk Management

### 8.1 NCRIS NCI Project

With the NCRIS NCI Project now implemented in its entirety, there are few residual risks for the program until its termination in 2011.

The following tables update the risk analysis for the NCRIS project and present it in a form consistent with that adopted for the EIF project (as shown in the EIF Project Plan).

#### 8.1.1 Adequacy of Resources

	<b>Risk</b>	<b>Mitigation Strategy</b>
1.1	<p>Sufficiency of co-investment</p> <p>Impact: Major Likelihood: Unlikely Risk Severity: Medium</p>	<p>The NCI (NCRIS) Project contained a high dependency on co-investment and, at the inception of the program, there were substantial concerns about the adequacy of the co-investment to provide internationally competitive infrastructure and quality services.</p> <p>Strong levels of co-investment have been achieved within the NCI Project and are continuing to grow—encoded within the NCI Partner Service Agreement. The assessment opposite reflects the residual risk within the NCRIS program, with the major risks associated with co-investment transferring into the EIF Climate HPC Facility (Sec. 8.2.1)</p>
1.2	<p>Currency Fluctuations</p> <p>Impact: Likelihood Risk Severity Low</p>	<p>The previous Business Plan highlighted risks associated with currency fluctuations relating to the purchase of the peak system in 2009.</p> <p>At this time there are no residual risks associated with currency fluctuations.</p>
1.3	<p>Sustaining ongoing operations of the Merit Allocation (Flagship and Open Access) Schemes</p> <p>Impact: Critical Likelihood: Possible Risk Severity: High</p>	<p>The conclusion of the NCRIS Funding Agreement and the associated funding stream implies that the recurrent costs associated with Merit Allocation Schemes (Flagship and Open Access) may not be able to be met beyond the end of 2012 (see previous discussion in Sec. 6.3.1). NCI seeks discussions with DIISR in the coming year to address this potentially damaging matter.</p>

#### 8.1.2 Services

	<b>Risk</b>	<b>Mitigation Strategy</b>
2.1	<p>Quality of Service</p> <p>Impact: Major Likelihood: Possible Risk Severity: High</p>	<p>Previous international reviews have made highly favourable assessments of National Facility services provided by ANU through APAC. More recently, the strong uptake in partnership continues to confirm the value and quality of service perceived by major research organisations.</p> <p>A survey of users is planned for this business year as an ongoing quality assurance measure.</p>



	<b>Risk</b>	<b>Mitigation Strategy</b>
2.2	<p>Mismatch in organisational cultures adversely affecting service delivery through the third-party organisations</p> <p>Impact: Major Likelihood: Possible Risk Severity: High</p>	<p>The Specialised Facilities and aspects of the Computational Tools and Techniques programs are delivered through third party organisations.</p> <p>The selection processes that have led to these arrangements has ensured a suitable concordance of culture between ANU and the partner organisation. The requirements of NCI are encoded in legal agreements which specify the obligations of the service provider(s), and provide for termination clauses in the event of default.</p>

### 8.1.3 Security

	<b>Risk</b>	<b>Mitigation Strategy</b>
3.1	<p>Sufficiency of funds for data centre and system procurements</p> <p>Impact: Critical Likelihood: Unlikely Risk Severity: Medium</p>	<p>The physical security is reviewed by ANU and the present level of security is seen as adequate.</p> <p>The physical assets are ensured for their full replacement value.</p> <p>User (/home and /project) data is backed up on tape on site and at a separate location. Users may request that (tape archived) data of significant value is backed up to tape at a separate location.</p>

### 8.1.4 NCRIS Objectives

	<b>Risk</b>	<b>Mitigation Strategy</b>
4.1	<p>NCI Project may not meet NCRIS objectives.</p> <p>Impact: Likelihood: Risk Severity: Low</p>	<p>The NCRIS objectives for the project (summarised as follows):</p> <ol style="list-style-type: none"> <li>1. Sustaining a capability HPC system commensurate with international norms,</li> <li>2. Advising DEST on investments relating to specialised systems,</li> <li>3. Operating a community based merit allocation scheme,</li> <li>4. Providing advisory and support expertise to research communities,</li> <li>5. Maintaining a supported strategic plan for national computational needs,</li> </ol> <p>either have been met, or are in the process of being met, with the exception of (5), for which the Commonwealth has indicated its preference for this matter to be considered by NRIC and AeRIC. In the case of (5), this is being pursued with more limited scope for the climate / earth systems science community through the objectives of the EIF project.</p>



## 8.2 EIF Climate HPC Project

### 8.2.1 Financial — Infrastructure and Sustainability

	Risk	Mitigation Strategy
1.1	<p>Sufficiency of funds for data centre and system procurements</p> <p>Impact: Major Likelihood: Possible Risk Severity: High</p>	<p>The indicative split of funds between the data centre (\$20-25M) and system procurements (\$25-30M) is seen to be adequate to meet the objectives of this project. This assertion is supported by a 2007 data centre design study undertaken by consultants engaged by ANU, a recent (2009) consultant report about purpose built data centres and associated infrastructure, and the recent and ongoing experience of technology developments in HPC, through relationships with vendors and other major data centres.</p> <p>In March 2009, ANU commissioned a design study by a firm of international data centre consultants to examine potential sites on the ANU campus and to specify the building design and requirements, given a brief that the structure may need to be extended in future years. This scalability, and in particular the need to procure only those services (mechanical and electrical) that are needed to provide for the 2012 systems, leads us to believe that the funds allocated to the data centre allocation will be sufficient. At the time of writing, we understand that a data hall of around 900-1,000 sq. m. and matching space for ancillary equipment will be achievable within the funds available. This will be confirmed in the consultant's report which will be completed in April 2010.</p> <p>The HPC system will be of a modular/scalable design and the funds available will be sufficient to acquire a substantial system. NCI's work to date in developing a high level system specification suggests that it should be possible to acquire a peak performance of around one petaflop in 2012 with the funds available. (The sustainable performance, however, will be somewhat less.) Of course a major determinant of whether this can be realised will be fluctuations in the Australian dollar against US currency. ANU will investigate mechanisms by which this can be managed.</p>
1.2	<p>Long-term operational sustainability</p> <p>Impact: Critical Likelihood: Possible Risk Severity: Extreme</p>	<p>NCI is strongly sustained by a range of partnerships (which are steadily building) with major national research organisations and consortia.</p> <p>Operations of the planned EIF funded facility will need to be sustained on a strong partnership base, the foundation of which will comprise the current and future major partners, ANU, CSIRO, and BoM. Informal discussions within the framework of the NCI Steering Committee confirm the willingness and intention of these organisations to sustain reasonable operational costs.</p>



	Risk	Mitigation Strategy
		<p>Negotiations to confirm co-investment involving major partners will commence formally in 2010Q2, with the relevant payment and expenditure milestones associated with this issue protecting the Commonwealth's investment.</p> <p>NCI is also increasing its outreach to support major research centres/collaborations, in keeping with its mission, and with a view to increasing co-investment from leading research organisations.</p>
1.3	<p>Additional expertise needed to develop the Centre</p> <p>Impact: Critical Likelihood: Possible Risk Severity: Extreme</p>	<p>The long-term success of the investment is critically dependent not only on the infrastructure component but on the range, extent and quality of the expertise needed to deliver services. This expertise is needed to realise goals in the priority area (climate/marine), and also in other research to be supported, and it is particularly necessary in realising the goal of upscaling the impact of high-end computation on Australian research.</p> <p>The need for increased expertise and support is recognised within NCI which will work with existing and new partners to build the co-investment needed to realise the potential of the new facilities. Current outreach programs are demonstrating advanced research outcomes generated by high performance and data intensive computation, with NCI using these in seeking to attract additional co-investment from leading research organisations and partnerships.</p> <p>In the event that sufficient co-investment to establish a Development Team of sufficient size is not realised, it is likely that the desired scientific impact of the new HPC infrastructure will be significantly eroded.</p> <p>The strategies for mitigating this risk are the same as those outlined in risk area 1.2.</p>
1.4	<p>Depreciation Costs</p> <p>Impact: Likelihood: Risk Severity: Low</p>	<p>Elements of the infrastructure that comprise the new HPC complex will be integrated systematically into ANU business operations. The NCI Steering Committee will recommend to ANU that an agreed amount (a fixed percentage per annum of the current replacement value of the building) be set aside to provide for an effective long-term risk control against the emergence of a future maintenance issues for the building fabric and building systems. These maintenance costs will be incorporated into the ongoing operational costs for the complex.</p> <p>No provision will be made for depreciation of the HPC infrastructure.</p>



### 8.2.2 Suitability for Purpose

	<b>Risk</b>	<b>Mitigation Strategy</b>
2.1	<p>System procurement does not meet the requirements of research communities</p> <p>Impact: Major Likelihood: Unlikely Risk Severity: Medium</p>	<p>NCI has substantial and sustained experience in meeting the requirements of diverse research communities. Through direct interaction with researchers and research communities, through representation of senior researchers on its merit allocation committee, and through detailed specialised planning for the petascale facility (through the Services and Infrastructure Planning Group), the consultation process will deliver a specification which will meet the requirements of diverse research communities.</p>
2.2	<p>Data centre design and capacity is not sufficiently flexible to handle future requirements</p> <p>Impact: Major Likelihood: Unlikely Risk Severity: Medium</p>	<p>The data centre design is being informed by consultation with specialist data centre design consultants, HPC vendors, and through developing relationships with major international HPC installations. ANU has commissioned two studies (in 2009Q3 and 2010Q1) from consultants to advise on data centre design. There also exists substantial expertise within NCI and in ANU Facilities and Services in upgrading services and infrastructure to meet current generation systems in the existing data centre.</p> <p>The design brief incorporates space and utility requirements for current and future HPC and data services. Multiple site options are being examined, with extendibility being a key requirement.</p> <p>The risk analysis will be updated when ANU reports to DIISR on the chosen site of the data centre, following a decision informed by the data centre consultant's report.</p>
2.3	<p>Adequacy of access arrangements</p> <p>Impact: Moderate Likelihood: Unlikely Risk Severity: Medium</p>	<p>NCI has substantial experience in providing for national access to its facilities and services, both on merit and through partner share. This contract will extend the current arrangements by overlaying "access on merit" with partner shares to provide priority access for climate/ESS, flagship access for high impact research from leading researchers and research communities, and open access to accommodate the breadth of computational research requirements (complementing project funding from national granting agencies with computational services from NCI). Priority access principles will be formulated by the Project Steering Committee and approved by the NCI Steering Committee.</p> <p>NCI further advocates the establishment of a National Merit Access Scheme to coordinate access to resources available on the national peak systems. It is noted that alternatives to national coordination have been tried previously in various settings internationally and have not succeeded, ultimately being replaced by nationally coordinated systems. The present US Teragrid allocation processes are a clear example.</p>



### 8.2.3 Construction

	<b>Risk</b>	<b>Mitigation Strategy</b>
3.1	<p>Delays in design and construction</p> <p>Impact: Critical Likelihood: Possible Risk Severity: High</p>	<p>Design studies are already underway within ANU, with assistance from consultants. The schedule provides for a comprehensive design phase of nine months during which the project scope and user requirements will be carefully defined. ANU has substantial experience in managing complex construction projects and has recent relevant experience in substantially upgrading the existing data centre (Leonard Huxley Building) to accommodate a contemporary system, and in delivering this, on time and on budget, through its external contractors.</p> <p>The construction timeframe (of 12 months) is tight but is achievable, based on recent experience (outlined above) which established new utility services (power and cooling) for the Leonard Huxley Data Centre in a purpose built structure. Project management by ANU and its contractors builds in contingencies to accommodate construction delays (e.g., bad weather, industrial action etc).</p> <p>The risk analysis will be updated when ANU reports to DIISR on the chosen site of the data centre, following a decision informed by the data centre consultant report.</p>
3.2	<p>Compliance with relevant environmental or heritage requirements</p> <p>Impact: Major Likelihood: Possible Risk Severity: High</p>	<p>ANU is required to implement plans and strategies to manage its heritage places in order to best protect heritage values, in keeping with the requirements of the Environment Protection and Biodiversity Conservation Act 1999. As well as the Commonwealth legislation, the University is also required to ensure that its heritage management is compliant with all relevant State and Territory heritage legislation. All projects are assessed (using University staff or specialist consultants) against the heritage and environmental criteria identified in the University's campus heritage strategy documentation. The results of these deliberations form a key component of the documentation submitted to the National Capital Authority (NCA) as part of the formal submission for project works approval.</p> <p>The construction timeframe is tight and any protracted heritage or environmental delays will impact on the project. The preferred sites are not considered to have any significant architectural, cultural, heritage or environmental significance.</p> <p>The risk analysis will be updated when ANU reports to DIISR on the chosen site of the data centre, following a decision informed by the data centre consultant report.</p>



	<b>Risk</b>	<b>Mitigation Strategy</b>
3.3	<p>Adequacy of project governance and management</p> <p>Impact: Major Likelihood: Unlikely Risk Severity: Medium</p>	<p>ANU will undertake, govern and manage this project on advice from the NCI Steering Committee which will establish a Project Steering Committee with representation of the key co-investing stakeholders, and chairs of three project planning and implementation groups which have carriage of various facets of the project.</p> <p>The governance and management framework for the data centre follows established ANU practice. The specification of the data centre requirements and the detailed design and construction are separated, with the former, representing NCI, acting as a client to the latter that will deliver the design and construction to meet these requirements.</p> <p>All capital works undertaken at ANU are governed within an integrated policy framework of project management, client representative arrangements, and quality control. The HPC Data Centre construction will be carried out under the ANU Project Management Method whereby the Project Manager manages and coordinates the work of the consultants and contractors to ensure a quality result for the University, viz., that the project is delivered within budget, delivered at an agreed target date, the work meets University standards and complies with BCA standards and statutory requirements, and the satisfaction of the client group.</p> <p>Responsibility for the delivery of the project lies with the ANU Facilities and Services Division and a senior officer will be appointed as a project coordinator to manage the day to day implementation of University procedures as well as overseeing any government reporting requirements. A commercial Project Manager will also be appointed to manage the delivery of the project and to oversee the associated administration, including monthly reports, financial reporting, project scheduling, and budget management.</p>
3.4	<p>Adequacy of any sub-contracting arrangements</p> <p>Impact: Major Likelihood: Unlikely Risk Severity: Medium</p>	<p>Over the past decade, ANU has pursued an active infrastructure development and campus renewal program, coordinated by the Facilities and Services Division. Management of major individual projects is outsourced under the direction of the Division to specific project managers who are very experienced in undertaking construction projects on the campus. The new facilities will be designed by leading architectural and services firms with experience in designing data centres. A key criterion will be that the facilities achieve exceptional performance in relation to energy and water usage, and greenhouse gas emissions.</p> <p>ANU will procure resources for project design, management and construction in accordance with existing proven practices. The ANU will engage a quality Project Manager for construction activities, as well as quality trades. The project will have an agreed procurement strategy, with a construction program that reflects appropriate lead times for key equipment, and allows for all necessary inspection and test points by consultants and the Project Manager.</p>



	<b>Risk</b>	<b>Mitigation Strategy</b>
		<p>The project program will be monitored and reviewed through monthly reports to the Project Control Group as required by Government. The Program will include historically consistent allowances for weather delays.</p> <p>The ANU only engages Project Managers and major trades that are compliant with the requirements of the Office of the Federal Safety Commission (OFSC).</p>
3.5	<p>Management of costs during construction</p> <p>Impact: Major Likelihood: Possible Risk Severity: Medium</p>	<p>During the design phase, ANU undertakes progressive checks of project costs against the approved budget to ensure there is no design creep above client priorities. Value management reviews are undertaken at specific hold points (Preliminary Sketch Plan and Final Sketch Plan stages) to align scope with budget and for specific issues when necessary. The ANU client will sign off the Final Sketch Plan before proceeding to final documentation and tendering to industry.</p> <p>Appropriate allowances are made in the budget to cover cost escalation over the duration of the project as well as appropriate contingency allowances for design and unforeseen construction issues, e.g., latent conditions. ANU also identifies, where appropriate, “tradeables” that can be excluded from the scope of work, or which can be added, if the budget allows.</p>

### 8.2.4 System Procurement

	<b>Risk</b>	<b>Mitigation Strategy</b>
4.1	<p>Adequacy of design specification</p> <p>Impact: Critical Likelihood: Unlikely Risk Severity: Medium</p>	<p>The specification of services and infrastructure to meet the needs of research communities commenced prior to the formal start of the project. This will be formalised through the establishment of the Services and Infrastructure Planning Group. This group will define the services, and determines the broad infrastructure requirements, which subsequently will be developed into specifications to be formalised in an open request for proposal.</p> <p>NCI/ANU has substantial expertise in high-end system specification and this, together with vendor relationships, and the prior planning of services and infrastructure, will ensure a design specification that is fit for purpose and which will meet the requirements of the user community.</p>
4.2	<p>Adequacy of expertise and experience to undertake the procurement process</p> <p>Impact: Critical Likelihood: Unlikely Risk Severity: Medium</p>	<p>The governance and management of the procurement largely mirrors the recent successful procurement by BoM and ANU of interoperable HPC systems.</p> <p>The technical expertise in the NCI National Facility at ANU and the Bureau of Meteorology is strong well regarded, nationally and internationally, and has a track record of several successful major installations.</p>



8.2.5 Services

	Risk	Mitigation Strategy
5.1	<p>Peak system procurement does not meet the requirements of research communities</p> <p>Impact: Critical Likelihood: Unlikely Risk Severity: Medium</p>	<p>NCI management and the National Facility support team have substantial experience in meeting the requirements of diverse research communities, evidenced most recently through the joint procurement by ANU and BoM of interoperable peak facilities.</p> <p>The design of the infrastructure and services for this project will result from substantial community engagement through the Services and Infrastructure Planning Group, and through the project governance framework in which all major stakeholder organisations are represented.</p>
5.2	<p>Services to be put in place by 2012 do not adequately support the priority area (climate change, earth systems science, and water management) or other key research fields.</p> <p>Impact: Critical Likelihood: Unlikely Risk Severity: Medium</p>	<p>Similar to the response for 5., the project governance framework, together with the expertise in NCI and stakeholder organisations, is sufficient to ensure the design of a comprehensive service layer to meet the requirements of the priority area and other key areas of Australian scientific research.</p> <p>The greatest concern relates to the adequacy of funds needed to augment the support team to put in place the expertise (i.e., the development team) needed to maximise the science impact from infrastructure being put in place. This is dealt with below, with the risk assessment here relating only to the adequacy of the service design process.</p>
5.2	<p>Adequacy of funds to provide for advanced computational services, particularly for additional staff expertise to develop the centre's impact</p> <p>Impact: Critical Likelihood: Possible-Likely Risk Severity: High-Extreme</p>	<p>The intended impact of the planned petascale facility will be greatly diminished in the absence of a support team which is augmented substantially from level of that presently supporting the work of the NCI National Facility.</p> <p>Additional expert computational scientists are needed for the NCI support/development team to assist researchers in increasing the scalability of existing codes, in designing and implementing new, highly scaling codes and data-intensive services, and in assisting the outreach needed to assist the Australian research community to take up advanced computational methods.</p> <p>In the EIF environment, this cost must be born through co-investment, with this Business Plan having outlined a process by which major stakeholders will agree to the sustenance of the operations of the centre. This process will also emphasise the importance of a support team that is commensurate with the infrastructure and the goals of this project.</p> <p>NCI will also seek to increase the level of co-investment provided by other major research organisations and consortia. NCI has a record of success in attracting co-investment and will build on this, with increasing emphasis on outreach activities, and by closely positioning NCI with the support of high-impact research communities from leading research organisations.</p>



	Risk	Mitigation Strategy
		<p>In 2012, it is estimated that the cost of each support EFT will be less than 5% of the recurrent cost of electricity. This evolution in the balance of operational costs provides a reason for optimism that the requisite complement of expertise can be sustained.</p>



## 9 Key Performance Indicators

### 9.1 Performance Indicators — NCRIS NCI Project

The performance indicators listed below are unchanged from those agreed in the initial NCI (NCRIS) Project Plan, and which have been reported on in the Annual Reports for the financial years 2008–09 and 2009-10.

Objectives	Key Performance Indicators	Outcomes
Develop a national strategy for advanced computing to support eResearch in Australia.	<p>A national advanced computing strategy to support eResearch is in place and maintained.</p> <p>Demonstrated progress towards the implementation of the strategy.</p>	A long-term commitment by NCI and its partners to strengthening the Australian advanced computing infrastructure for eResearch in line with international trends.
Strengthen relationships to provide national cooperation on Australia's advanced computing infrastructure.	<p>Membership and participation in NCI categorised by degree of involvement and contribution of organisations.</p> <p>Funding received from sources other than NCRIS.</p>	NCI leading the development and use of an Australia-wide advanced computing infrastructure for eResearch.
Improve the peak computing capabilities of the NCI NF to serve the demands of Australian researchers.	<p><b>Achievements</b> Key research achievements obtained by users of the NF.</p> <p>Significant developments in the computational services and tools on the NCI and partner facilities.</p>	Significant research achievements by users of the NCI NF.
	<p><b>Capability</b> Capability of the NF relative to countries of similar size and development such as Canada, Sweden and Korea.</p> <p>Capability of the NF relative to the needs of the Australian research community as indicated by extent and type of demand.</p>	Australian advanced computing capabilities commensurate with comparable countries.
	<p><b>Accessibility</b> Extent of the use of the NF indicated by the number of states represented by users, projects, users and resource allocations.</p> <p>Availability and performance of the NF indicated by the available system units, used system units, efficiency of operation and extent of parallel computation on the NF.</p>	Australia as a leader in the development and use of high-end computing services for eResearch.



Objectives	Key Performance Indicators	Outcomes
	<b>Customer Service</b> Summary of surveys of users of the NF showing their level of satisfaction.	Demands of major Australian researchers for advanced computing services satisfied by the NCI NF.
Implement specialised computing capabilities through NCI SFs to serve the demands of Australian researchers.	<b>Achievements</b> Key research achievements obtained by users of SFs.	Significant research achievements by users of the NCI SFs.
	<b>Capability</b> Capability of SFs relative to the needs of the Australian research community as indicated by extent and type of demand.	Australian advanced computing capabilities commensurate with comparable countries.
	<b>Accessibility</b> Extent of use of SFs indicated by the number of States represented by users, projects, users and resource allocations.  Availability and performance of Specialised Facilities is comparable to the NF.	Demands of major Australian researchers for advanced computing services satisfied by NCI SFs.
	<b>Customer Service</b> Summary of surveys of users of SFs showing their level of satisfaction.	Demands of major Australian researchers for advanced computing services satisfied by NCI SFs.
Develop and deliver outreach services for users of advanced computing infrastructure.  Participation and contribution of organisations in the Outreach program.  Number and types of outreach activities conducted.	Extent and nature of working linkages with research organisations.  Increased skills of users of advanced computing infrastructure.  Extent of investment by partner organisations.	Increased uptake of advanced computing and in research organisations.  Increased levels of co-investment in advanced computing.



## 9.2 Performance Indicators — EIF Climate HPC Project

Objectives	Key Performance Indicators	Outcomes
Quality of the infrastructure and value for money	Capability, reusability and extendibility of the Data Centre	<ul style="list-style-type: none"> <li>• Suitability as a general purpose, high-end data centre</li> <li>• Capability for future reuse, expansion and repurposing of the data centre fabric and ancillary services.</li> </ul>
	Capability of the Petascale HPC System	<ul style="list-style-type: none"> <li>• Infrastructure specified and procured, and services designed in collaboration with climate / earth systems science / water management and other research communities, that meet the intended requirements.</li> <li>• Capability of the peak system relative to international norms, and in particular the relativity to countries of similar size and development.</li> <li>• Favourable price / performance comparisons</li> <li>• Energy efficiency— approaching best practice, commensurate with the role of the “climate change” role of the facility.</li> <li>• Well balanced system performance — in accordance with world’s best practice.</li> </ul>



Objectives	Key Performance Indicators	Outcomes
<p>Quality of services and their acceptance by the user community</p>	<p>Uptake of services in the priority, flagship and open/general access areas — measured in terms of numbers of users with funded research, user communities, organisations, major research centres, etc that are supported</p>	<ul style="list-style-type: none"> <li>• Australia as a leader in the development and use of high-end computing services in research.</li> <li>• Outcomes in priority area (climate/ ESS /water) acknowledged by major stakeholders in project.</li> <li>• Outreach activities increase research engagement with high-impact research centres/communities.</li> </ul>
	<p>Demonstrable impact of NCI services on the research and practice of supported communities, particularly in the priority access areas and the flagship projects</p>	<ul style="list-style-type: none"> <li>• Increased high-impact research outcomes (according to the usual norms) supported by the NCI National Facility.</li> <li>• Demonstrable outcomes of high impact that are of national benefit (to be defined with major stakeholders).</li> <li>• Significant improvements in the uptake and use of highly scaling codes on the peak system.</li> </ul>
	<p>Extension of services to high profile research communities which are new or non-traditional users of high-end computing services</p>	<ul style="list-style-type: none"> <li>• Growth in the uptake of significant services from groups/communities which are new to high-end computing.</li> <li>• Increased numbers of projects and users associated with these projects.</li> </ul>
<p>Strength of the sustaining partnership and collaborations</p>	<p>Development of a long-term sustaining partnership</p>	<ul style="list-style-type: none"> <li>• Signing of a collaborative agreement (milestone) sufficient to provide for recurrent operations.</li> <li>• Commitment by sustaining partners to the development of long term strategic and operational planning in facilities and services to support climate change, earth systems science, national water management.</li> </ul>
	<p>Cash contributions and commitments sufficient to sustain and develop operations without unacceptable compromise</p>	<ul style="list-style-type: none"> <li>• Sufficient contributions to expand the computational science expertise base needed to advance the scope and depth of the services / applications.</li> </ul>



Objectives	Key Performance Indicators	Outcomes
	<p>Development and delivery of outreach services which increase uptake and co-investment opportunities</p>	<ul style="list-style-type: none"> <li>• Outreach activities developed which build research capability and which increase the uptake of high-end computing services.</li> <li>• Increased skills levels in the research community.</li> <li>• Increased levels of co-investment arising from increase interest and uptake.</li> </ul>
<p>Research impact of the new petascale infrastructure</p>	<p>Infrastructure provides a substantial and demonstrable improvement to Australia's science research capability</p>	<ul style="list-style-type: none"> <li>• New science challenges addressed and outcomes arising, particularly in the priority and flagship areas.</li> <li>• Measurable increases in research output and research impact in key fields.</li> <li>• Substantial national benefit outcomes (to be defined) in the priority area.</li> <li>• Support for and uptake by national high profile research collaborations, e.g., ARC Centres of Excellence</li> <li>• Substantial benefits to the nation through new services in climate / ESS / water management (presumes sufficient co-investment to provide for the expert staff to develop these services).</li> </ul>
	<p>Demonstrable upscaling in the sophistication and ambition of Australian computational science, transitioning to levels evident in international peer centres</p> <p>(Dependent on sufficient co-investment to put in place a strong software / code scaling development team.)</p>	<ul style="list-style-type: none"> <li>• Significant changes in the profile of usage of National Facility with increased use of highly scaling codes.</li> <li>• Increased impact and national benefit of outcomes, through more ambitious science questions being asked / answered.</li> <li>• Ultimately, greater uptake of, and co-investment in, high-end computing as the substantial benefits of high-end simulation are realised in the research community.</li> </ul>



Objectives	Key Performance Indicators	Outcomes
External benchmarking and review	Establishment of substantive relationships with international peer centres (and related organisations)	<ul style="list-style-type: none"> <li>• Establishment of a number of significant relationships which assist NCI in its long term planning and development.</li> <li>• Potentially, some collaborative activities in areas of common interest.</li> </ul>
	User surveys of the quality of service	<ul style="list-style-type: none"> <li>• Evidence that the services provided by the NCI National Facility and its Support Team are meeting the needs of the communities served, and with a high level of quality.</li> </ul>
	(Possible) External review of the computational service and its research impact	<ul style="list-style-type: none"> <li>• Evidence that the services are meeting the needs of the communities served and are of high quality.</li> <li>• Validation of the planning framework and directions.</li> <li>• Insights into new directions and practice based on international perspectives.</li> </ul>