

**NCRIS PLATFORMS FOR COLLABORATION
NATIONAL COMPUTATIONAL INFRASTRUCTURE COMPONENT
INTERIM IMPLEMENTATION PLAN
(PROJECT PLAN)**

1. Description of the Project including a full description of the entire infrastructure that is to be funded under this Agreement

Project Name: The National Computational Infrastructure (NCI) Component of the Platforms for Collaboration or the National Computational Infrastructure Project in short, hereafter known as ‘the Program’.

The investments in the Program will deliver an internationally significant ‘tier one’ computing capability which can be assigned on a merit and priority basis, build essential expertise in High Performance Computing (HPC) needed to support NCRIS-supported and other priority research and provide a national strategy for computation infrastructure. The investments will build on and progress work to date on the development of compute infrastructure, especially relating to capabilities beyond the reach of departmental or even institutional investment levels and which make sense to share and re-allocate over time.

OBJECTIVES

The Program:

- sustains an Australian capability computing system commensurate with international standing and developments;
- advises DEST on further investments in the national computing fabric relating to ‘specialised systems’ that can preferentially support application oriented ‘tier one’ computing requirements;
- operates a community based merit allocation scheme that grants access to the NCRIS funded resources; and
- provides advisory and support expertise to research communities benefiting from NCRIS computational infrastructure investments; and
- maintains a supported strategic plan for national computational needs.

The Program has three activities, organised into sub-programs as follows:

Activity 1: Planning, Access, Policy

- Strategic **planning, including** development of business and marketing plans;
- **Access** to facilities and services through a merit and priority allocation process;
- **Policy** Development, including acceptable use policies,

Activity 2: Operational Facilities and Services

- the **National Facility** with a peak computing system for ‘capability’ computing;
- the **Specialised Facilities** that provide high-end computing systems that complement and strengthen the National Facility in its support for user communities and their applications; and
- **Computational Tools and Techniques** activities that assist users to develop advanced computing tools and applications on these facilities.

Activity 3: Outreach

- A coordinated set of **Outreach** activities to facilitate the uptake of advanced computing.

GOVERNANCE & MANAGEMENT

The Program is to be governed by the ANU on advice from a steering committee (referred to herein as the NCI Steering Committee (NCI/SC)), and managed by a program director (referred to herein as the NCI Director) with the support of a program management office (referred to herein as the NCI Office).

The ANU will appoint a Program Delegate to whom the NCI Director will report, and through whom the Chair of the NCI/SC will transmit reports to the University. Hereafter references to ANU are to be taken to be actions undertaken through or by the Program Delegate on behalf of the ANU.

The Program is expected to have significant relationships with a number of organisations over time and of varying nature and purpose. Of particular importance are NCI Partners, which are those organisations operating or providing significant resources (either funding, infrastructure or other in-kind resources) or services within the Program. The NCI/SC will determine the resource level required to be considered a partner; an average annual contribution of \$1M may be a reasonable starting point. NCI Partners will be acceptable as such to DEST.

NCI Steering Committee

The NCI Steering Committee (NCI/SC) will advise the ANU on:

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

The NCI/SC will make these recommendations based on considerations and reports on all the NCI activities.

The NCI/SC is directly representational, and will consist of:

- an independent Chair;
- one representative from each of the NCI Partners;
- one representative from each of the NCI Prospective Partners;
- the Chair of the Merit Allocation Committee;
- the NCI Director (non-voting);
- one representative of NCI affiliates.

Provision will be made for alternates on the NCI/SC. The NCI/SC Chair and alternate will be appointed by ANU, and will be acceptable to all NCI Partners and DEST. Resolutions of the committee on concerns of national importance or on significant risks to meeting objectives of

the Program will be incorporated in Reports to the Government. The Chair has the casting vote if required.

The current members of the NCI/SC are:

- Independent Chair: Professor Mark Wainwright, AM
- Chair NCI Merit Allocation Committee: Associate Professor Brian Yates
- NCI Director: Professor Jim Williams (acting)
- Partner: Professor Robin Stanton (ANU)
- Prospective Partner: Dr Alex Zelinsky (CSIRO)
- Prospective Partner: Dr Neville Smith (Bureau of Meteorology)
- Prospective Partner: Dr Chris Pigram (Geoscience Australia)
- Affiliate: Professor Doug McEachern (representing Go8 DVC-R's)

NCI Director and Office

The activities of the NCI component will be managed by the NCI Director supported by the NCI Office.

The NCI Director will be appointed by the ANU, on advice from the NCI/SC and will be acceptable to DEST. The Director is accountable to the ANU for the overall co-ordination and integration of the Program according to the principles set out in section one of the NCRIS Roadmap.

The NCI Director with the support of the NCI Office will:

- coordinate the development and management of operational plans and other strategic planning documents and submissions, specifically the Annual Business Plans;
- manage the Program in accordance with those plans and as directed by the ANU (acting on advice of the NCI/SC);
- develop and manage Program related agreements and relationships with third party organisations including NCI Partners;
- develop annual budget and financial statements;
- coordinate progress reports from subcontractors; and
- provide the NCI/SC with regular performance reports.

The NCI Office will be established within the ANU under arrangements whereby ANU will provide administrative services including human resources management services, financial services, contract and legal services and information technology support services. Such services will be provided on a cost recovery basis.

The NCI Office will provide the NCI/SC's secretariat.

Merit Allocation Committee

A Merit Allocation Committee (MAC) will independently evaluate applications for resource grants on the systems funded by the Program.

The ANU will appoint the membership of the MAC, on advice from the NCI/SC. The MAC will include senior researchers who are able to review the applications competently. In addition to the general membership of the MAC, ANU will appoint a Chair of the MAC, with the Chair to be acceptable to DEST. Provision will be made for DEST to appoint a member

or observer to the MAC. The current APAC MAC is determined from nominees of the APAC Partners. NCI will review the MAC within the implementation phase.

The MAC will report to the NCI/SC.

The role of the MAC will be to make allocations granting access to the NCRIS share of the resources on the National Facility and other funded facilities. The principles of the Merit Allocation Scheme by which merit allocation will be determined will be outlined in the Annual Business Plan and agreed with DEST. For grant allocations in 2008, the members of the APAC MAC will continue and the APAC MAC assessment criteria will be used.

User Forum

The NCI Director will convene and chair a NCI User Forum. This forum will enable the NCI/SC and NCI Director to better meet their responsibilities and help inform their activities. This forum will be complemented by specialised workshops with focus groups, including those that may be convened from time to time by DEST.

Forums will, where possible and where appropriate, coincide with major outreach activities.

Coordination with NCRIS Capabilities

ANU recognises the role of the Program within NCRIS, and the relationship between the Program, other NCRIS capabilities and other components of the Platforms for Collaboration capability (Pfc). The envisaged role of AeRIC will also be taken into account. The NCI Program will contribute to the coordination of NCRIS using the following principles:

- will give priority to the NCRIS roadmap for the capability areas;
- actively participate in the ICI and ANDS policy settings;
- preferentially utilise services developed under the ICI and ANDS projects.

Mechanisms for cooperation and consultation will be developed at all levels, both with regards the other Pfc components and more broadly with the NCRIS capabilities.

PROGRAM ACTIVITIES

Planning, Access and Policy Activity

The initial Program funding for these and the outreach activities is \$2M, and this funding also includes the management activities. Given the nature of the activities and sub-programs set out below, further development of the Program budget is expected to highlight the need for additional funding in this area. This additional funding requirement within the Program budget, with \$1M being probable, will be addressed in subsequent business plans.

Planning Sub-Program

The ongoing relevance of the NCI capability, within the NCRIS roadmap, depends on processes that consider the activities of the NCI, and their continued development. Therefore strategic planning will take place to consider the broader development of the activities within NCI. This planning will require input from the user communities, and will address how facilities and operations areas develop, including the relevance and strengths of specialised systems with respect to the major investment in the National Facility.

The role of the National Facility is covered in the section on National Facility.

The key criteria to establish a specialised facility (SF) system will be:

- the extent of the demand for the facility for the 'high-end' users in the Australian research community and particularly the NCRIS areas established in the planning process;

- the access to the SF and user support that will be provided for researchers Australia-wide;
- the proposed investment in the SF represents good value in comparison to other possible investments;
- the host organisation is capable of providing the required level of service to users as a component of the ‘national advanced computing infrastructure’; and
- The host organisation is able to implement policies, system scheduling and processes to accommodate projects granted resources under the Merit Allocation Scheme.

NCI commitment to an SF system will not be made without DEST approval and will be based on advice of the NCI/SC. Emphasis will be on application focussed systems and facilities within the NCRIS capabilities and other areas of national interest such as climate change. Initial approaches to NCRIS capabilities 5.1, 5.3, 5.10, 5.12 and 5.13; and to ANU, CSIRO and BoM will be extended to determine their interest in such facilities. Based on feedback from these discussions, a more refined and formal call for proposals will be considered.

Alternative approaches, such as an open call with criteria as above, and emphasising requirements around ‘national advanced computing infrastructure’ will also be considered by the NCI/SC. Budgets for this sub-program will be proposed by the NCI/SC for inclusion in annual business plans.

Planning will also include appropriate marketing of the NCI capabilities, both in the capability areas and the broader user community. This may include stronger representation in the international settings that help establish Australia as an attractive place for high-achieving researchers.

Access Sub-Program

This sub-program will review various merit and priority access processes, with a view to determining the most appropriate process for the NCI. Areas highlighted for development include multi-year allocations, further changes to data allocations in light of ANDS, coordination of merit processes across peak and specialised facilities, the composition of the merit committee, and shares allocations.

Further, the sub-program will also examine changes in shares depending on the needs of the community, and purchase options for research and non-research users. The sub-program will also re-evaluate access conditions as part of any technology change occurring within the Operational Facilities and Services Activities.

Policy Sub-Program

A range of policies will be required to meet the goals of the Program. These may include acceptable use policies as technology developments change; policies resulting from access considerations; and operational considerations of the NCI facilities.

Outreach Activity

Outreach relates to the engagement of the broader community in the development and use of national computational infrastructure and services by focussing on the following strategies:

- increasing the awareness of the capabilities and benefits of advanced computing;
- increasing the generic skills of users and potential users of advanced computing; and
- promoting the use of advanced computing.

Outreach is expected to build on the community of interest in advanced computing and successful outreach events that APAC and its partners have conducted. These include Conferences and Exhibitions, 'all hands' meetings, summer schools and internships.

Activities may include workshops focussed on technologies and applications relating to the interests of advanced computing such as virtual reality and visualisation systems, and representation at relevant national and international events, such as supercomputer and data management conferences in the US and Europe.

The APAC07 Conference in October will continue the successful APAC conference series with over 200 delegates and workshops on relevant technologies.

NCI will have a presence at SC07, the largest international supercomputing conference.

A website <http://www.apac.edu.au/nci> has been established to provide information for users and other interested parties about the Program.

Operational Facilities and Services Activities

The National Facility (NF) Sub-Program

The NF will provide world-class advanced computing services to Australian researchers by providing access to computing and data systems to satisfy the requirements for 'capability' computing.

The ANU will continue with arrangements to assist continuity for users between the services provided under the APAC Agreement and this Agreement, taking advantage of the overlap in services being offered under the two agreements. The call for applications under the Merit Allocation Scheme will be issued in early October 2007. In accord with NCRIS principles, grants will be openly available to any researcher from any organisation that is eligible to seek access to capabilities supported by NCRIS, and based on merit.

The NF will operate under a 'resource share' model whereby the level of contributions accepted from organisations determines the resources on the facilities that are allocated to each organisation for their users.

The arrangements will allow external organisations to access fixed resource allocations (specified number of processor hours) as well as shares of the system depending on the nature of their investment and interest.

A share of available resources will be made for large and long term commitments, and a fixed resource for smaller and shorter commitments. The competing timeframes with machine lifetimes, annual budgets, and research project requirements is a subject of policy formation in NCI.

The NF will also offer 'start-up' grants to help researchers evaluate the computational and data requirements of their research and to prepare applications for the Merit Allocation Scheme.

Peak Computing Capability

The NF will plan the installation of enhanced peak capability at times mutually agreed between DEST and ANU. Between the scheduled peak system installations, the NF may also replace or enhance its development system in order to explore new technologies and to provide additional computational capacity and meet alternative configuration requirements. The development system should support a useful subset of work and provide improved computational capacity to the peak in a cost effective way.

The business case for an initial development system costing around \$1M will be developed with procurement occurring in early 2008 if approved by the NCI/SC.

All significant purchases involving NCRIS funds will be through open competitive tender.

Data System and Services

The NF may also operate associated data systems, services, specialised software, and provide associated expertise to support data collections and data intensive applications.

The system is intended to provide a range of data services including:

- domain specific and general data analysis and mining software;
- specific file-type manipulation software;
- other software packages, tools and libraries and porting additional packages;
- data storage by fast index relational database, flat file and directory;
- application support for various data transfer methods, data streaming and data access;
- methods of managing access rights and authentication; and
- data usage, audit and performance profiling and instrumentation tools.

Grants to use the associated data systems to support data intensive research will be provided through the Merit Allocation Scheme.

ANU will sustain the existing data collections support provided by APAC through the funding provided by this agreement for the NF through to 30 June 2008. Further resource provision and merit allocations will be made through processes to be established by the Australian National Data Service.

Due to ANDS not yet being ready to provide services, NCI will, with DEST approval, issue a MAS Data call to support data collections in 2008.

The NF will manage a process whereby storage resources, above and beyond the resources provided by the facility under the merit process, can be allocated to NF users.

System and Application Software

The NF will continue to provide and support software including application packages, support tools, compilers, databases and file manipulation tools and parallel programming libraries, performance tools and debuggers. The software will be regularly updated and associated tools installed as improvements are sought by users.

Common Datasets

The NF will also manage significant datasets that are shown to be useful for multiple projects of the facilities or that provide a broader service. The need for these datasets will be reviewed with respect to the needs of the users in a manner similar to other services, software and tools managed by the facility.

Training, User Forums and Workshops

The facility will provide training material and host workshops on making effective use of the facility, and improving skills through advanced techniques and tools provided on the system.

User forums will be organised to provide user input to the developments of the facility and promote development of strategic user communities using the facility.

Physical environment

It is envisaged that ANU will provide the physical environment for the computing systems of the NF including computer room space, environmental services, facilities management and necessary facility upgrades to accommodate new systems.

The manner in which ANU might also provide access to existing ANU infrastructure, including the data system and visualisation facilities, network connectivity and communications and software licences will be included in the Annual Business Plans.

For the Implementation phase, ANU will provide access to its infrastructure relevant to NCI in the same manner as occurred under APAC.

Specialised Facilities (SF) Sub-Program

Access to SF to support specialised high-end computation services will be arranged through subcontract agreements. The need for the SF will be investigated through the strategic planning activity of NCI including the consultation with the user community. Recommendations for SF will then be considered by the NCI/SC.

SF will provide to the Merit Allocation Committee details of resources to be made available to users. This will help inform the MAC in developing arrangements for users to gain access to these systems.

Through the NCI, SF may also operate a shares based model to allow other investments to take place through the NCI Program, similar to that of the National Facility.

Hosts of SF will provide the NCI Office with adequate accounting and reporting for the resources funded. In addition, these hosts will provide the NCI Office with appropriate audited accounts relating to the funding received, and report on and assist reviews of the performance and reliability of the facility and its services.

Computational Tools and Techniques (CT&T) Sub-Program

Activities within CT&T aim to improve the software development environments for users of the facilities in NCI. It is anticipated that there will be an increasing need to provide researchers with problem-solving environments which expose the methods to solve the problem and hide the complexities of the systems. It is also expected that some projects will contribute to international efforts by adapting codes and supporting them on the NF and specialised facilities.

The NCI/SC will establish processes that select and support a small number of projects to provide enhanced computational tools and techniques to targeted user groups. The projects may be based around any of the computing resources or services provided by any of the organisations with a relationship with NCI.

The criteria for supporting a project will include:

- The extent to which it improves computational research for a targeted user community
- Improved effective use of the facilities by the general user base on these facilities;
- The extent of the user demand for the tools and techniques; and
- The track record of the key people in the project team and their relationship with user groups.
- Support for the cohesiveness of the national computing fabric.

2. Details of the access and pricing arrangements to ensure that the infrastructure is accessible by researchers on the basis of merit.

Access to systems and services supported by the Program will be by merit determined through a merit allocation process that operates independently of the organisations providing the services and is open to any researcher from any organisation that is eligible to seek access to capabilities supported by NCRIS.

For the NF, an organisation may also purchase a portion of the systems for their priority use at full operating cost. The cost to research organisations will be the annual cost of the facility in that year factored by the share made available. The cost of access to non-research and commercial users will probably be set at a higher rate, depending on the impact of market forces.

For SF, the NCI/SC will establish a process by which proposals can be formed around the establishment of application specific computing installations. Proposals will only proceed to implementation with endorsement from DEST. For those facilities, an organisation may also purchase a portion of the systems for their priority use at full operating cost. The cost to research organisations will be the annual cost of the facility in that year factored by the share made available. Again, the cost of access to non-research and commercial users will probably be set at a higher rate.

For CT&T, an open call will be made for project proposals that satisfy the criteria of the sub-project. The NCI/SC will establish an open review and decision making process for this purpose. A factor in merit for such activities will be the resources made available from other parties to the project, and the demonstrated ability to pursue the project as a single activity combining resources from multiple sources.

3. Details of the ownership and management arrangements that will result in efficient and effective operation of the infrastructure.

ANU will hold the systems in the NF purchased with NCRIS funds on behalf of and for the use of NCRIS and broader research community.

No physical assets are envisaged under the CT&T activities. In general, Intellectual Property generated by these projects will rest with the parties participating in the projects. Improvements to any pre-existing IP will rest with the owner of that IP. Specific project proposals may identify other arrangements between the parties.

The organisations hosting and responsible for the SF will own the assets and will provide services to users determined by the Merit Allocation Committee commensurate with the contributions from NCRIS.

4. A business case including full financial projections that will ensure the effective ongoing management of the infrastructure.

The Project is an integrated part of the PFC investment plan.

The provision of capability computing is a strategic decision which all comparable nations have pursued. While the overall envelope of funding that could be provided is a matter for priority assessment; there is a minimum funding level below which a capability solution is not feasible and a central system becomes unattractive from a research community perspective.

The four year financial projections for NCI are shown below, with the 2007-2008 projections highlighted:

Statement of Income and Expenditure July 2007 to June 2011

	1/7/07 to 31/12/07 (\$'000)	1/1/08 to 30/06/08 (\$'000)	1/7/08 to 31/12/08 (\$'000)	1/1/09 to 30/06/09 (\$'000)	1/7/09 to 31/12/09 (\$'000)	1/1/10 to 30/06/10 (\$'000)	1/7/10 to 31/12/10 (\$'000)	1/1/11 to 30/06/11 (\$'000)	Total (\$'000)
INCOME									
NCI Income	1000.0	3000.0	7000.0	0.0	9500.0	0.0	5500.0	0.0	26000.0
ANU	0.0	1000.0	1000.0	0.0	1000.0	0.0	1000.0	0.0	4000.0
CSIRO	0.0	1000.0	1000.0	1000.0	1000.0	1000.0	1000.0	1000.0	7000.0
APAC contributions	1527.0	753.2	0.0	0.0	0.0	0.0	0.0	0.0	2280.2
Affiliate cash commitments	0.0	60.0	20.0	0.0	0.0	0.0	0.0	0.0	80.0
ANU inkind	1200.0	1200.0	1200.0	1200.0	1200.0	1200.0	1200.0	1200.0	9600.0
Other Income	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
TOTAL INCOME	3727.0	7013.2	10220.0	2200.0	12700.0	2200.0	8700.0	2200.0	48960.2
EXPENDITURE									
<u>Operational Facilities and Services</u>									
<u>National Facility</u>									
- Systems and Upgrades	0.0	800.0	0.0	3000.0	9000.0	0.0	2000.0	1000.0	15800.0
Power and cooling infrastructure		300.0		4000.0				300.0	4600.0
- Facilities Management - computer room	290.0	290.0	290.0	325.0	325.0	325.0	325.0	325.0	2495.0
Facilities Management - data	370.0	370.0	370.0	370.0	370.0	370.0	370.0	370.0	2960.0
Consumable power	200.0	200.0	300.0	600.0	700.0	900.0	900.0	1000.0	4800.0
Software and network	150.0	150.0	150.0	150.0	150.0	150.0	150.0	150.0	1200.0
System and User Support	1200.0	1200.0	1200.0	1200.0	1200.0	1200.0	1200.0	1200.0	9600.0
- Other	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	2210.0	3310.0	2310.0	9645.0	11745.0	2945.0	4945.0	4345.0	41455.0
Specialised Facilities			4000.0						4000.0
<u>CT&I</u>		125.0	250.0	250.0	250.0	250.0	250.0	250.0	1625.0
<u>Planning, policy, access and outreach</u>	0.0	250.0	250.0	250.0	250.0	250.0	250.0	250.0	1750.0
TOTAL EXPENDITURE	2210.0	3685.0	6810.0	10145.0	12245.0	3445.0	5445.0	4845.0	48830.0
INCOME LESS EXPENDITURE	1517.0	3328.2	3410.0	-7945.0	455.0	-1245.0	3255.0	-2645.0	
Carry Over	0.0	1517.0	4845.2	8255.2	310.2	765.2	-479.8	2775.2	
FUNDS AVAILABLE	1517.0	4845.2	8255.2	310.2	765.2	-479.8	2775.2	130.2	

Notes

1. APAC has forward paid operating expenditure and system maintenance for 2008.
2. Minimal co-investment has been formally committed, and there is currently little expectation of additional significant sources.
3. Inkind and other co-investment and expenditure for Specialised Facilities has not been included due to ongoing uncertainties around these facilities, especially in the context of ongoing BoM/CSIRO discussions around an earth systems research system.
4. At the time the notional budget for the Project was developed, insufficient details on the cost of support plant and infrastructure and utilities were known. Since then, preparation of the 4 year plan for the National Facility has shown significant increases in these costs. In particular, the available budget only allowed for one peak system in

the NCI funding timeframe. The revised plan calls for a developmental system in 2008, with a new peak system in 2009 and an upgrade in 2010. In 2011 a new development system is proposed with the expectation of a peak system in 2012 funded outside the current NCRIS arrangements.

5. The expenditure in the proposed budget is based on a number of principles:
 - a. replacing the peak computing system in the NF with expenditure commensurate with previous purchases of peak systems;
 - b. obtaining around 50% of the resources in the NF for the merit allocation purposes;
 - c. investing in SF to obtain a significant resource share of these facilities for the merit allocation purposes;
 - d. supporting the CT&T activities to a level commensurate with the APAC CT&T program until 30 June 2007; and
 - e. ensuring expenditure for the NCI Office (management) is commensurate with its anticipated responsibilities.

The Steering Committee made the following observations at its December 13 meeting (taken from the accepted minutes):

All members agreed that there is not enough money, even if pooled to get what is needed.

Australia is losing international research competitiveness due to prolonged under-funding of computational infrastructure, and this situation is not sustainable. As an example, regional scale assessments for Australia as part of the IPCC 5th Assessment may not occur unless the situation is addressed.

Members agreed with the following general strategy

- *Do the minimum required to maintain some measure of international research competitiveness;*
- *A single peak system in 2009, expending essentially all NCRIS and likely co-invested funds;*
- *A new machine in 2011-2012 would require funding from other sources, and obtaining this funding would be an important role for the NCI Director;*
- *There are significant machine room and power expenses above what was identified in the notional budget.*

It was recognised that this was an expedient solution significantly constrained by budget, and that there was a strategy to address the situation.

There are distinct possibilities for ANU, BoM and CSIRO to work together, especially around the next research machine at BoM and with joint data centre options.