

PART THREE RESEARCH INFRASTRUCTURE

Investment Criterion 1 **An investment plan must result in excellent research infrastructure that addresses the national requirements of the relevant capability area described in the NCRIS Roadmap**

“Platforms for collaboration include the following sets of inter-related components:

- *Data storage management, access, discovery and curation* to improve interaction and collaboration;
- *Grid enabled technologies and infrastructure* to enable seamless access to the facilities and services required in various research fields;
- *Support skills* to assist researchers in developing and using this infrastructure effectively;
- *High performance computing* to allow analysis, modelling and simulation; and
- *High quality network access through high capacity bandwidth* to permit interaction with diverse data and computing resources.”⁴

While the consultations for 5.16 and the other NCRIS capabilities confirm the importance of these components, the importance of authentication and system wide authorisation has become clearer as has the importance of tools as the primary means for broadening the adoption of e-Research methodologies. Therefore support for *Tools* and *Authorisation* have also been planned as part of the investment.

The investment plan also takes into account the following sector characteristics:

- Institutions and federal and state agencies within the sector will spend internally on these same components, and in aggregate spend far more than NCRIS
- The inter-relatedness of components means that priority cannot be addressed by selecting some needs over others, for instance compute over data (or vice versa)
- The state of readiness of solutions, technologies, and associated products and services, is highly variable when viewed on a community or discipline basis

⁴ National Collaborative Research Infrastructure Strategy, DEST, February 2006.

At present, national general purpose e-Research infrastructure services are supplied by AARNet (expending circa \$40M pa) and APAC (expending circa \$20M pa). Outside of those organisations, the following can be noted:

- The APAC state partners spend about another \$25M pa on additional activities, many of which relate to Platforms for Collaboration
- The Bureau of Meteorology and CSIRO jointly invest more than \$10M pa in HPC and data services within the High Performance Computing and Communication Centre, and invest in additional related infrastructure, which needs to inter-operate with investments by NCRIS
- Data interests, such as Geoscience Australia, state and federal research agencies, and the humanities overall, need to be included within a common national e-Research infrastructure, in a co-ordinated approach to national data management as proposed to PMSEIC
- NCRIS capabilities themselves will spend an estimated \$10M pa on e-Research infrastructure within the various capabilities
- LIEF grants are in excess of \$30M pa with some component in e-Research infrastructure, and Universities themselves invest in e-Research infrastructure

Fortunately several DEST initiated Research Information Infrastructure projects have explored and demonstrated e-Research capabilities over the last few years; so that much is known about the nature of the issues and the means for addressing them.

Overall, awareness of the need for co-ordination and the importance of standards is growing. However, the policy barriers to collaboration and the impact of uncoordinated investment are less well understood.

Finally, the aggregate infrastructure expenditure mentioned above exceeds \$100M pa, and the aggregate ICT budget of research institutions is larger again. NCRIS investments need to add value to these investments and provide a framework in which they can inter-operate and be leveraged beyond their immediate intentions.

Vision

The vision laid out by the ERRC report has been consistently reinforced throughout the consultation process.

“Australian researchers will be world leaders in the use of advanced and innovative information and communications technologies, to achieve internationally recognised, high quality research outcomes across Australia's national research priorities;

Australian researchers will be able to access data, instruments, computing capability and to collaborate with each other, through advanced ICT, enabling them to engage readily in collaborative research and contribute significantly to the solution of major national and international research challenges;

Australian researchers will have the necessary education, training and skills, and support from ICT and information management specialists, to use advanced ICT efficiently and effectively;

The implementation of e-Research capabilities across the Australian research sector will provide a leading influence on the uptake and enhancement of such technologies by Australian business and industry; and

The Australian community and economy will benefit from the advanced capability enabled by e-Research”.

While some of these characteristics cannot be addressed by infrastructure development, the broad thrust certainly can.

Therefore: this Investment Plan builds on current activities to enhance, develop and deliver the infrastructure services needed to realise this e-Research vision.

Emergent Behaviours

National and international scale ICT systems are some of the most complex system built and operated and have behaviours and consequences that are difficult to plan or even predict. The idea of ‘emergent behaviour’ has been developed to represent the fact that IT development leads to changing social and economic behaviour because of complex interactions of multiple developments rather than because of individual technologies or functions. For instance, that mobile phones have changed the way

social interactions are planned and organised producing a more dynamic and fluid day-to-day life style, which was not foreseen when the technology was developed.

An interesting manner in which to set goals in e-Research is to describe desirable emergent behaviours and develop activities that can lead to those behaviours by deploying ICT technological capabilities in infrastructure development:

Area	Projected emergent behaviour
Data	All research data is suitably curated and retained, and co-ordinated services support easy location, access to, and analysis of, that data.
Grid	‘Problem oriented’ virtual infrastructures are routinely and easily constructed from any number of sensor, instrument, compute, data or visualisation resource, located anywhere.
Support Skills	A cohort of experts co-operatively operates and extends the infrastructure and readily assists users solve research problems.
HPC	An inter-operating infrastructure of peak, shoulder, institutional and departmental resources, matches capabilities to needs and priorities.
Networking	An inclusive, high speed low latency toll free network, connects all researchers to all research resources, including international sites.
Authorisation	Researchers control who can do what with their resources through a common method based on a single login for each researcher.
Tools	Users are able to work with familiar tools despite accessing multiple and remote compute, data, analysis, sensor or instrument resources.

e-Research Architecture

While crystal ball gazing in ICT is notoriously difficult, a medium term view of the above can be refined into architectural implications for e-Research infrastructure.

Possible achievements In the medium term:

Routine retention of research data is in place at custodial institutions, shared use of high capacity regional data centres is the norm, 'google like' collection meta-services operate across all published Australian primary research data

Multiple capability class computing facilities and associated infrastructure is in place, as are investments for regional shoulder and capacity class systems

Commoditisation of 'ensemble style' compute capacity is achieved so that common user tools can run anywhere

Co-operating centres of expertise in e-Research technologies, and co-operating centres of e-Research infrastructure delivery, operate and interact globally

Many 'problem specific' infrastructures are in every day use, integrating multi-state resources

An AAF is in routine use across the majority of research activities

Campus infrastructure is harmonised with a functional AREN with appropriate reach and bandwidth

Associated policy frameworks and best practice standards are published

Longer term goals

Networking and access control needs to become a non-issue for research activities, with dedicated virtual networks supplied on demand, connecting large collaborations of researchers, and allowing data and analysis to be shared within high quality virtual presence delivered from theatres to desktops.

Exemplar information services will need to provide access to knowledge through browser like functionality that front-ends advanced analysis tools and the automatic access and aggregation of appropriate data to answer user questions and provide 'expert advice'.

System level science and international research collaborations will need to be in continuous 24x7 operation drawing on the interaction of multiple tools and data source from across disciplines to address problems at a national and global scale; built by merely re-configuring common underlying services.

National e-Research Architecture

Data will trend towards larger data centres located on a national, regional or sub-regional basis, because raw retention is a 'lowest common denominator problem' with commodity solutions, economies of scale, and little if any business differentiation; and the need for co-location of computing and the delivery of ultra-high bandwidth integration will drive costs up in smaller sites.

High end computing will be trend towards locating with large scale data, to support timely analysis over those holdings, and to support the information products that will become the most effective, and sometimes the only effective, means of providing access for many users.

Therefore the e-Research architecture will be a network of computing and data resources of various scales, coalescing over time; so that private computing exists with private data; as the data scales up, the computing will scale up; and as the data is shared, the computing will also be shared.

A second tier of large scale data sources will include aggregators of highly distributed data, large scale in-silico research facilities, and high volume image generators in physical research facilities.

As a consequence, the 'network' will separate into two layers, as the 'system-like' demand diverges from the 'human-like' demand. The human layer will be largely web based from researcher to researchers and services. The high performance infrastructure layer will be system to system and able to divert bandwidth to extreme scales of burst transactions, commit multi-point bandwidth for sustained periods to high levels of QoS, and interact with transaction schedulers.

The research network will also become geographically diverse to better match the data demand, and reach from researcher to service, and the capacity and QoS needed system to system.

Above this foundation infrastructure, a distributed layer of standardised middleware services will support workflows, and allow migration of analysis to data rather than of data to analysis.

The national architecture will then become one of inter-operating information services, where access to those information services implies access to data, computing and network resources, a middleware layer supporting workflows and migration of analysis to data, international agreements on content formats and meta data standards, and a shared expression of rights, roles and identities.

The architecture will require national services to assure integrity over the shared expression of rights, roles and identities by authoritatively aggregating and re-purposing appropriate information.

However, the unique mix of access and provenance that might be defined by the participants, sponsors and custodians for each data set leads to devolved content management. Hence the architecture must differentiate between content and systems; and a national policy framework and accreditation for shared services and service providers will be required.

Platforms for Collaboration

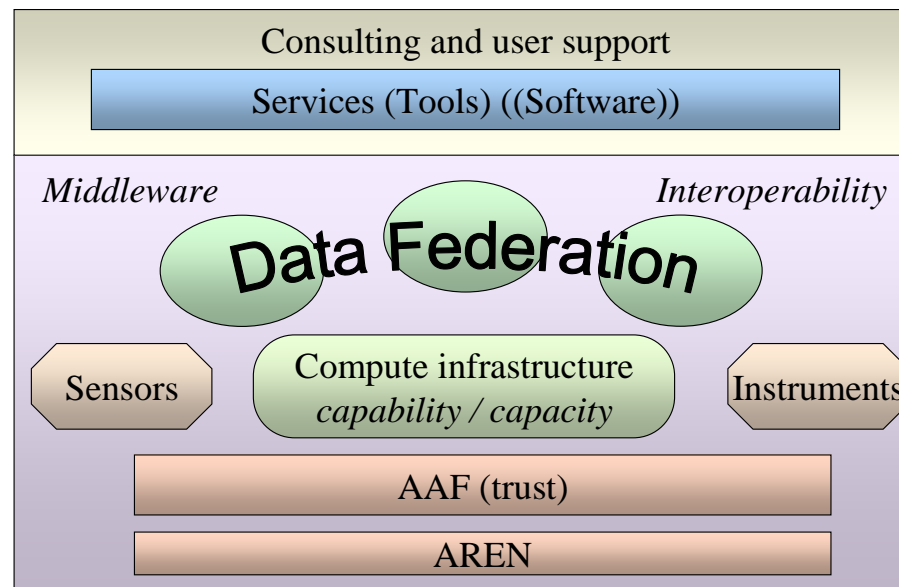
Given these over-arching directions, and noting the gathering dominance of data, the investment plan has developed around a view (depicted to the right), that given fundamental capabilities in access (networking and authorisation services) and given the provision of compute infrastructure for processing and modelling, and given data gathering from sensors and instruments; that the broader collaborative e-Research requirement can be best described as needs around data: for data access, data sharing, data analysis and data re-use. (See section seven for confirming international views).

It is in the communal use of data within and across multiple disciplines that e-Research opens the most valuable doors. From an infrastructure point of view, this requires a large variety of tools commonly in use across widely different disciplines to all be able to access and use the same data, and for researchers to be easily able to control who can do what with which elements of their data.

The NCRIS consultations have also strongly endorsed the view, that beyond access services for data, access by researchers to suitable tools and user support is a critical factor in enabling e-Research. Essentially, effective tools and information services are the only means by which e-Research can be made feasible for most researchers; and of value beyond the current cohort of the “user-builders” of cyber-infrastructure.

Consequently, the investment plan focuses on a number of components of the overall situation:

- The continued provisioning of the AREN and its use and access policies
- The provisioning of trust services that enable collaborative pan-organisational access control (the Australian Access Federation, AAF)
- The continuation and 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 (ANCF)
- The creation of the Australian National Data Service (ANDS) that can provide key capabilities, including federation services, outreach activities and stewardship of some classes of data, as well as source the expertise and provide the leadership needed to crystallise data capabilities within research organisations
- The provisioning of a national grid environment that brings relevant tools with associated data and compute capabilities to researchers (interoperation)
- Support for nationally integrated user support operating across all services and all service providers



A set of basic capability components were developed with various communities and then a number of alternative arrangements were considered. The final framing of the investment has focussed around four ‘programs’.

- A program to advance and harden the next generation middleware components needed for e-Research in selected disciplines and priority research areas
- A program on data infrastructure, to establish a national data collections service and to provide the expertise and outreach activity needed by the sector
- A program on computational infrastructure, in the form of a peak general purpose computational system and additional classes of systems
- A program on inter-operation and collaboration services, to support system inter-working, builder’s tools, and user services

A small component of funding will also be applied to support infrastructure in the Australian Research and Education Network and the Australian Access Federation, as well as supporting the NRICC committee and the eRI program office.

Expertise

One of the most important issue for an effective e-Research infrastructure concerns the development and provision of appropriate expertise.

Platforms for collaboration itself and nearly all the other investments in e-Research infrastructure across the NCRIS capabilities will increase demand for expertise.

The investment plans show that the various communities are at different stages of development towards an e-Research perspective; which means they will necessarily have access to very different levels of such expertise. This will be true more broadly.

Also, e-Research involves the use of multiple and entirely unrelated specialisations, such as curation of data, advanced networking, or parallel software for supercomputing. Added to this, grid capabilities and middleware are a rapidly evolving set of specialisations in their own right (such as searching, authentication and authorisation). Research groups cannot possibly cover this space.

Some important factors related to e-Research expertise are as follows.

- Expertise development and access can be enhanced by building groups of specialists rather than relying on unrelated individuals who can easily become overloaded and unresponsive
- Different expertise and different levels of expertise are required during different stages of a communities migration toward e-Research
- Eventually some expertise needs to be embedded in communities (eg. data curation) and some needs to be embedded in service providers (eg. network management)
- Along the way, flexible collaborative teams are needed so that the infrastructure can evolve as the requirements are better understood

Also, as communities become more e-Research oriented, they tend to co-evolve services for data generating and gathering, with services for information analysis and re-use. This happens because each community needs to develop a consensus on the standards required for inter-operation, and that consensus is largely driven by practical experience.

So e-Research depends on standards and standards development is always a long iterative processes. This means that researchers will necessarily undertake bespoke software development in order to continue their research while their communities converge on discipline specific standards.

Further difficulties then arise as reliance is placed on such software leading to a need for improved software engineering and particularly software productisation expertise.

The time required to develop highly generic standards, such as for general data management or job description, are even longer so that discipline specific middleware and tools can be expected to proliferate for some time to come.

Many of these issues were identified by the e-Research Co-ordinating Committee and the basic perspective developed by that committee remains valid and is reflected here.

A general solution to expertise is beyond the budget of Platforms for Collaboration as is a broadly based education campaign is most likely out of scope from an infrastructure investment point of view.

However the problems arising from the fragmentation of e-Research infrastructure is exacerbated by this missing expertise and in particular the leadership that could be expected from high levels of expertise.

Therefore an approach to expertise development, its internationalisation, and the means to access that expertise, needs to be part of the business plan.

Specifically, each program should identify expertise related to system or data interoperability, build expert groups at an international level of quality, support access to that expertise in an advisory mode, contribute experts to development projects on an agreed basis, and demonstrate the ability to participate in and influence global development of standards, components and middleware in that area.

The goal will be to identify and develop identified pools of national expertise in the important aspect of the middleware and service offerings created by Platforms for Collaboration, particularly those relating to middleware and inter-operability.

Program: The National Data Infrastructure

This program area will bring together data management interests to work towards the vision:

All research data is suitably curated and retained, and co-ordinated services support easy location, access to, and analysis of, that data.

The activities in this program are intended as a step towards a visionary and landmark research infrastructure that could meet Australia's research data needs.

To create momentum in that direction, the program will help experts and the builders and users of curation, access and analysis tools; to work together to improve community data management practice; to 'keep data better'. It will provide:

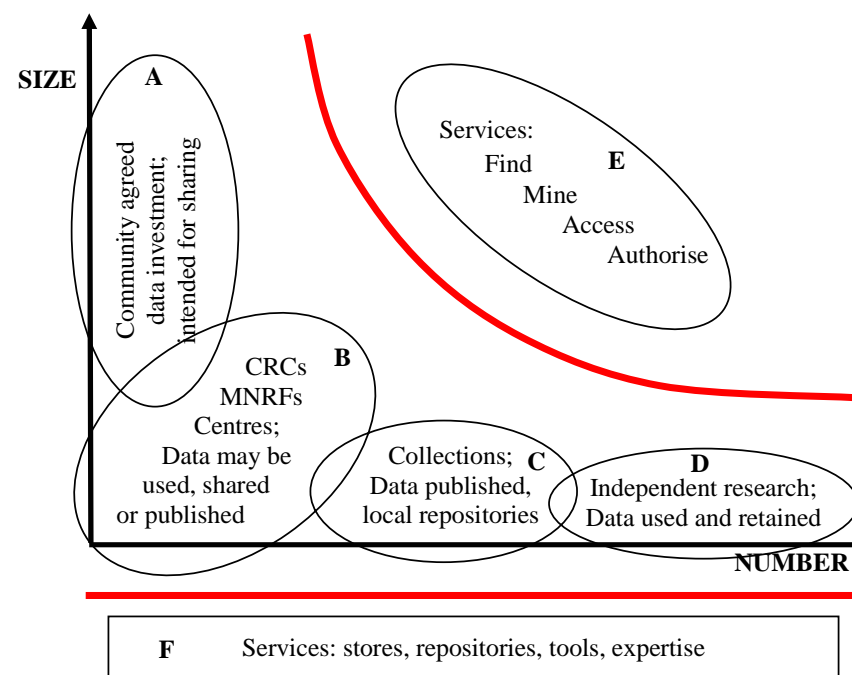
- Outreach advisory and support services to assist researchers establish data curation practices
- Outreach advisory and support services to assist institutions establish research data management infrastructure and practice in a nationally coherent manner
- Federation level services, including data location, access and data mining
- Data management brokering to help match needs to providers
- Stewardship services hosting some categories of collections
- A community based merit allocation scheme for managed resources

Consultations and surveys undertaken by Platforms for Collaboration suggest that data volumes are growing rapidly and that much research data is held in isolated forms and not easily accessible.

Generally, two kinds of data can be distinguished. Private data is often held by a researcher through self-generation, or by acquisition from surveys, an instrument or device. Public data refers to data placed in a public repository for general access.

Often, researchers also make some private data available on a controlled access basis, but do not normally provide public access to all the data they hold. Also, where data is derived from significant public or shared investment, a common data holding protocol often applies where researchers are granted controlled access to the data collected on their behalf for a limited time. After that time the ownership of the data reverts to the researcher's host institution, or is made public through a repository.

With reference to the diagram to the right, in developing a rationale for NCRIS investment, the following missions, properties and responsibilities seem to be present.



- (A) Several communities capture and share data of general value, where large datasets are generated or gathered as the result of significant investments; such as in astronomy, high energy physics, earth observations and bioinformatics. The size of these data sets tends to be very large, typically in terms of tens to hundreds of terabytes and sometimes petabytes, with high growth rates.
- (B) More frequently, organisations, such as CRCs and MNRFs, have data gathering, generation, curation and publication as part of their mission; or use data for internal purposes, the results of which are published as an information services; or they provide research support services that generate data. The data often is privately held, and only processed information is published. The size of these data sets tends to be in the many gigabytes to terabyte range.
- (C) Even more frequently, many research organisations, departments, teams, and even individuals have established data collections the contents of which are intended for publication and access. These publicly available data sets

(donated by researchers) are often kept in institutional or personal archives or web sites. The size of these data sets tends to be in the multiple gigabyte range.

- (D) Nearly all researchers generate or store data on the desktop, much of which is only ever intended for individual use and which is inaccessible to other researchers, visible only by publication of derivative research results. Most of this raw data is intrinsically hard to re-use and could only be accessed if it was moved off the desktop. The size of such data sets is highly variable.
- (E) Some organisations provide access to research data, either within disciplines or more broadly, and may or may not hold the primary data themselves. These organisations provide catalogues, abstracts or thumbnails, and can search and possibly mine data over which they hold indexes. These sources can only easily provide access to public data. The aggregate size of data accessible in this way would be in the range of terabytes to petabytes.
- (F) A final class of mission relates to the retention, curation, access and analysis of data sets. Generic service providers necessarily operate only on public data, although some specialised services could operate on controlled access data. The co-location of data sets and the consequent ease of datamining and linking is often part of the value of these services. Many institutions might operate such a service on their own behalf for their retained private data. The size of repositories tends to be in the range of hundreds of terabytes to many petabytes, and exabytes would be ultimately conceivable in some settings.

As the discussion suggests, data is everywhere, and the aggregate investment across all of the missions and user communities is large and rising quickly, and the importance of retaining, cataloguing, valuing and re-using data is unquestioned.

It also shows that data management services will be provided by a variety of sources and investment by NCRIS needs to be based on principles that fit within a broader framework. The first two columns in the table below are derived from policy under development by the AVCC, NH&MRC and ARC and the third summarises the investments developed here.

The investment proposal was also informed by PMSEIC, which noted:

“The Data for Science Working Group discussed at length the idea of a new National Centre for Data for Science. There was considerable support within the Group for a Centre; it was felt that such an initiative would be of benefit and may be a useful mechanism for progressing many of the above recommendations.

The Working Group considers that there is a range of functions that a Centre could assist with, including:

- *Facilitating and promoting the changes reflected in the recommendations;*
- *Working with those in specialist scientific disciplines to discover datasets;*
- *Establishing vital repositories; and*
- *Working collaboratively with the research, government and business communities to support the proposed new approaches to data for science.*

The Working Group stopped short of recommending the establishment of a Centre. The Working Group concluded that the high-level expert committee (recommendation 2) should decide whether such a centre was desirable and, if so, where it may be hosted, and what its role and governance mechanisms should be.”

<p>Institutions should:</p> <ul style="list-style-type: none"> ▪ Develop and implement a policy on data ownership ▪ Provide guidelines to researchers on ownership, what to keep and researcher responsibilities ▪ Maintain durable records on what research data has been held and ensure that research data is under the control of the institution where the work was performed ▪ Provide secure systems for holding data and for granting access to that data 	<p>Researchers should:</p> <ul style="list-style-type: none"> ▪ Determine what data to keep, considering research community practice and any project or legal requirements ▪ Ensure research data is retained (for at least 5 years from publication of results) using institutionally provided mechanisms ▪ Ensure at the end of employment (for whatever reason) data retention passes to the institution ▪ Maintain confidentiality where it exists 	<p>NCRIS should:</p> <ul style="list-style-type: none"> ▪ Provide national implementations of the federation services described under categories E and F ▪ Build expertise and provide outreach services that can assist others, including ‘training the trainers’ ▪ Broker solutions for collections and researcher needs and support identified collections to fast track e-Research development where appropriate ▪ Ensure promulgation and use of agreed legal frameworks and templates governing access to data
---	--	--

The Australian National Data Service (ANDS)

The Platforms for Collaboration investment will establish an Australian National Data Service (ANDS), as a cooperative centre of expertise in research data management.

ANDS will address the needs raised by PMSEIC by providing three mutually reinforcing services focussing on shared services, stewardship, and outreach.

While it would be possible to focus on any one of these and still provide value, the development of a national centre of expertise will be significantly enhanced by embodying and bringing together knowledge from all three.

With reference to the previous analysis, of the two meta-service categories depicted in the diagram, those related to re-use (E) and those related to retention (F), data search, data mining and access control services in (E) and national data location and movement services in (F) are systemic, beyond the bounds of individual institutions, and complementary and value adding to individual data holdings.

ANDS will therefore focus on providing common services in support of research data collections and to provide integration infrastructure that facilitates sharing of data, so that researchers can more easily discover, access, use, analyse, and combine digital resources as part of their activities.

While simple services will come first, the longer term vision is to connect Australian and international data centres, repositories and online collections to enable serendipitous discovery, cross-disciplinary research, and cross-repository workflows and to transform the disparate collections of research data around Australia into a cohesive corpus of research resources.

ANDS will also provide specific services that support collections of public data (starting with collections in category (B) and (C)) that could be preserved and made accessible. This service adds considerable value to the data and will not otherwise be easily or quickly provided. Over time, specific services may grow or diminish as research institutions establish a position on research data management.

ANDS will also be heavily involved in the identification, installation and adoption of user-centric tools, and the engagement of the research community and stakeholders through consultation and outreach activities. ANDS personnel involved in outreach will also be grounded in day-to-day data management and data services, (and vice-versa).

In general, ANDS will not set up new infrastructure facilities, but will extend and build-on existing or proposed capabilities.

Unlike the national compute infrastructure which has the benefit of six years of investment in APAC, the development of ANDS will need to proceed more slowly, so that potential participants can understand the proposed functionality and negotiate their involvement.

NRICC will:

- Establish an inaugural Steering Committee to provide guidance on the further development of ANDS (in consultation with DEST)
- Invite expressions of interest in ANDS from potential participants (in consultation with the Steering Committee)
- Develop the form of the ANDS agreement and initial program (in consultation with DEST and the steering committee)
- Finalise the program plan and execute the program agreement (in consultation with the Steering Committee and proposed participants)
- Appoint a Director: Australian National Data Service, who will also be the Data Collections Program Manager within eRI

ANDS will:

- Receive funds from NRICC under the arrangements defined in the program agreement and against the purposes and activities identified in the program plan
- Operate as a single co-operative centre across a network of contributing partners allocating aggregate resources against needs through an agreed merit process
- Engage the Steering Committee and others in developing a community of interest around research data management, engage in relevant international developments
- Develop and sustain the operation of federation services that enhance access to research data
- Develop and sustain an outreach activity that assists researchers and research organisations to improve their research data management practice
- Develop relationships with custodial service providers to deliver an effective national collections management service
- Manage resources provided to ANDS within the terms of the program agreement
- Report to NRICC on the implementation of the program plan on a regular basis
- Propose a revised project plan on an annual basis

Technical Services

The establishment of an increasingly sophisticated set of data access services over federated repositories is the primary goal of ANDS. While those services will be defined by ANDS, they will be implemented by appropriate service providers.

In the short term, the following priority services and infrastructure elements have been identified:

- A registry of research collections, including search and access services to those collections
- Persistent identifier services for digital research objects
- Data re-positioning services between repositories for performance, policy and privacy reasons
- A metadata schema registry to encourage standards and assist interoperability of data

Expertise around the services already exists in Australia, and the intention is to adopt, adapt and influence global systems and standards, engaging in international co-development where appropriate.

Once these federation support services are established, common data analysis and visualisation services, generic data quality assurance services, notification and data curation services, and common data submission and presentation services, will be targeted.

Future activities would encompass data fusion services, cross-discipline terminology services, aggregated statistics service, workflow services, etc.

ANDS will assist NRICC identify data management tools that could be developed in other investments to improve data management at the desktop.

Outreach Services

This activity aims to improve the expertise, capabilities, and most importantly, organisational policy and everyday practice related to research data production, management, structuring, description, analysis, and curation.

The intention is to support researchers, data scientists, and IT infrastructure staff understand the developing requirements around research data management in order to participate in this new age of digital collections; and to assist managers gain confidence in suitable policy options they might adopt.

As a result, the ANDS outreach will help develop a national research data management framework.

ANDS will actively create networks of expertise to bring together researchers and research technologists to develop recommendations against data management issues. It will:

- identify a strategic agenda around skills development and awareness building for research data in Australia and convene appropriate events, workshops and seminars.
- provide access to regional sources of advice to support individuals and projects develop improved data management practice.
- enhance this outreach by proposing activities within the discipline services program to develop data architectures for specific communities and multi-disciplinary research projects.
- liaise with other government data initiatives, such as the ABS National data network, and water, geospatial and health data networks.
- collect and disseminate information on best practice in e-Research data management policy

Stewardship Services

The intention of the Stewardship activity is to contribute to the development of solutions for long-lived public collections of data.

It will comprise a federation of organizations with an intrinsic interest in data management. NCRIS funds will extend those services in ways that meet the ANDS objectives.

Participation in the stewardship activity will be open to research data centres in universities and national institutions such as CSIRO, the National Library and the National Archives.

ANDS will focus on providing co-location and access methods for community-level collections which occupy the middle-ground between national reference collections and individual researcher collections. On-going support for such collections is often beyond the reach of a single institution and below the threshold that could attract recurrent support at a national level.

Pre-requisites for a data collection to be supported by ANDS would include: a high level of community engagement, high quality structured data, adherence to international standards, and the on-going participation in data curation activities by the community.

ANDS will fund hosting services in relevant data centres, and expertise in data quality assurance, data manipulation, and data management. Curation will be a shared responsibility between the ANDS and the identified research communities.

Budget

The likely budget scenarios are under development.

Stewardship

A level of funding will be provided, increasing over time, through which ANDS will be able to assign support for collections.

The detailed estimates are difficult, as different data and different access regimes demand different service support levels and associated costs.

For instance, some collections will need some curation work as part of their acceptance, some data will require information services to be hosted, some data will require higher access rates and a larger proportion of underlying disk to tape support, and some data will require off-site replication and so on.

The approach is that the funding will remain relatively set and the capacity available for collections will depend on the requirements of the holdings as they are accepted.

The rising funding line and the trend in storage system capabilities will ensure that a significant number of new collections can be accepted each year.

ANDS will have a range of capabilities in its retention providers so that the needs for different collections can be met by appropriately configured and costed services.

ANDS will also provide a brokerage function to assist other collections identify providers most related to their area of interest.

Federation services

The scale of cost associated with the federation data services is more related to their development and extension over time than their annual operation.

While this could be provided in activities in other programs, it is placed in ANDS to ensure that governance arrangements around ANDS can determine the priority for the tools and services to be developed.

As ANDS becomes more operational, this funding level will decline, so that the more advanced location and analysis services are more user community determined.

The ANDS program will however support the installation and operation of servers and the configuration and maintenance of the services when operational.

Outreach

The outreach program will be funded on an in-kind co-investment basis with organisations that also seek to address the curation difficulties of their local research communities.

Some support will be provided in every region.

ANDS will work to identify communities to which curation assistance should be provided for improving curation or data management practices, and to also identify collections to be accepted into the stewardship activity.

The program agreement will provide that a significant portion of this effort is to work on collections identified for stewardship by ANDS.

Program: The National Computing Infrastructure

This program area will bring together co-investors in major compute infrastructure and related service providers and work towards the vision:

An inter-operating infrastructure of peak, shoulder, institutional and departmental resources, matches capabilities to needs and priorities.

The program will evaluate and review investments in computational infrastructure that support services which can best be funded and justified at the national level.

A strong objective is to ensure the requirements of Australia's advanced computational science can be met, especially through the provision of peak computing and mass data storage systems.

In addition, the program is intended to support the more purpose specific computing capabilities needed by other NCRIS investments.

The capabilities will be extensible and allow 3rd parties to co-invest for dedicated capacity or to 'purchase' capacity dedicated to discipline or problem specific services. The intention is to maximise options for economies of scale, to build-on rather than compete for expertise, and to provide a well managed environment able to support computing capacity and services developed as part of other NCRIS investments.

The APAC National Facility has operated over the last six years supporting the high end, peak or capability class of user, with a quality of service and benefit to research that has been regularly reviewed. Each review, including the most recent in 2006, has praised the facility, its management and customer support and commented that it provides an exemplary contribution to Australian national research.

Countries comparable to Australia are planning the installation of additional major computing infrastructure to support leading-edge research. At the high end, the US National Science Foundation is envisaging a 'petascale computing' environment that would see at least one system with a peak performance in the 1-10 petaflops range by 2010. It is likely however that a petascale system will be installed before that date.

However, several important comments on HPC were made during the APAC review:

- The peak facility should be retained at or near the historical level
- A more frequent purchase rate is needed to improve the return in Tflops delivered against dollars, assuming that a significant overlap in the operational periods of systems is manageable within the machine room infrastructure

- Given the already competitive nature of access, more resource is likely to be needed to allow for the broader clientele envisaged under NCRIS

Estimates of likely demand suggest a ten-fold increase in compute and data storage capacity every 3-4 years. A reasonable broadly stated goal would be to provide Australia researchers with access to at least one system in the 100-200 teraflops scale by 2009-10 and a petascale system by 2012-13.

The program builds around the proposals submitted by APAC for the APAC National Facility to its 2006 review and commended by that international review panel.

Therefore the program proposes to extend the concept of the APAC National Facility to allow for additional systems at other sites and to allow for evolution in its implementation, such as more frequent purchase of overlapping (but somewhat smaller) systems.

Examples of specialised systems that could be provided as a national service need to be reviewed in the context of the program, but include:

- software 'engines' for commonly used applications software (e.g. Blast)
- a 'rendering' engine for visualisation at remote sites
- a content management system for streaming content such as animations (storing video and multi-media data, capable of streaming media to users)
- a bioinformatics data server
- analysis support in characterisation, particularly at the synchrotron

The consultations in the NCRIS context around tools and the processes by which e-Research activity can be more broadly adopted suggest the inclusion of a generic computing environment for commodity tools and broader e-Research services.

The Australian National Computing Facility (ANCF)

The ANCF will be comprised of four activities:

- A steering committee comprising members of NRICC and others to develop an integrative framework for aligning significant compute infrastructure requirements and investments
- A national peak capability service commensurate with international practice focussed on advanced services for the cohort of expert users with research goals that can only be met by this capability
- A national computational fabric supporting expert users and infrastructure builders, and hosting e-Research computing services
- A community based merit allocation system to consider system priorities and allocate resources against research community needs

Unlike, ANDS, the investment plan is able to build the development of the ANCF around the functions and management arrangements for the existing APAC National Facility at ANU.

Therefore ANU will continue to host the peak computing systems.

It would be possible to consider opening the hosting function up to contenders, however this is of no value to NCRIS unless a second source of funds were to be supplied to the merit allocated peak capability as a result.

The current strategy of acquiring peak systems every 3-4 years will be revised to acquire a peak system every 2 years and to operate peak systems with a long overlap (1-3 years). The length of the overlap will be determined by factors such as the needs for an overlapping service, floor space and energy as well as maintenance costs.

The national capability facility will also provide systems, specialised software and associated expertise to support data intensive applications. The data may be generated from the computational systems or be part of a data collection hosted by the ANDS.

The facility will continue to provide and support software consistent with having unique and specialised system. This software includes 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.

It is already the practice that the cohort of expert managers of high end facilities informally co-operate on tenders, evaluations and purchase arrangements for each major purchase. The ANCF will be a body of expertise in HPC and its requirements and solutions, and provide access to that expertise to assist other HPC centres.

In addition to the capability facility, ‘shoulder’ and specialists systems will be supported to provide a more flexible and purpose specific computing infrastructure.

This strategy will result in other sites being funded to provide specialised systems and skills relevant to priority research needs. It will also mean that the peak computing system can be focussed more on ‘capability’ computing and a more diverse range of systems and support will be available in response to user needs.

The basis for funding and the implementation arrangements will be based on the current arrangements applying to SII funding directed to the APAC National Facility.

The program agreement will ensure that 3rd parties are able to make arrangements to support extensions of these systems for the use of specific research organisations.

NRICC will:

- Building on the APAC foundation, establish a Steering Committee to oversight the program (in consultation with DEST) and confirm the interest of participants
- Finalise the program agreement and program plan (in consultation with the steering committee and participants)
- Provide funds to activities identified within the strategic plan, under terms defined in the program agreement
- Appoint a Compute Infrastructure program manager who will chair the steering committee and manage activities as agreed within the operational plan

The Compute Infrastructure program will:

- Undertake regular reviews of the needs of key user groups, determine the most appropriate upgrades to systems to meet these needs and recommend the allocation of funds to support those systems
- Continue the existing APAC merit allocation scheme and extend it to grant resources on all facilities funded as part of the program through a cooperative decision process that matches user needs and available systems.

National Peak Facility

The facility will be the e-Research expertise centre in high performance computing and will be expected to provide an outreach and consulting service.

The facility will be a continuation of the APAC arrangements.

Following the recommendation of the APAC review that the peak system should be replenished more frequently (with slightly smaller systems) and an overlap of operation planned, the intention is to shorten the purchase cycle to 2 and half years.

The recommendation is for an \$12.5 million acquisition at that time, which would be expected to yield a system in the 40-50 ranking on the TOP5 00 list.

The specific financial planning for the peak system is highly detailed and depends on expected purchase points, cash management between NCRIS and ANU, and costs associated with considerations around operational overlap.

Also, current facility management charges are being paid on a schedule agreed with ANU, facility management is currently forward paid to around September 2008 and a further \$1m will be required for the current peak system to the end of its maintenance contract (September 2009).

The program agreement will provide arrangements that permit 3rd parties to expand the peak system by purchasing a share or a fixed processor allocation, where the proposed use aligns with the research support intentions of NCRIS funding.

National Computing Fabric

The national computing fabric will provide the underlying computing capability to best support key disciplines and priority research of the other NCRIS investments.

The details of this investment are less clear but will include co-investment in systems and system support from the Program Members.

The overall financial plan is under development.

The program agreement will provide that half the system will be available for allocation under the common compute infrastructure MAS, and the remainder will be made available to support identified capabilities and priority research areas.

Capacity unused by either of those means will be available for the use of the Program Members at the sites where it is installed, as they see fit.

The program agreement will provide arrangements that permit 3rd parties to expand any of these systems through contract with the Program Member at which system is sited, where the proposed use aligns with the research support intentions of NCRIS funding. Such extra capacity can be allocated on whatever basis is agreed between the Program Member and its service customer.

Program: Interoperation and Collaboration Services

This program area represents the operational arm of the integrative infrastructure underpinning Platforms for Collaboration.

The program will bring together providers of shared compute and data services to operate an extended “Australian National Grid”, working towards the vision:

‘Problem oriented’ virtual infrastructures are routinely constructed from sensor, instrument, compute, data and visualisation resources

This emergent behaviour is most likely to arise through several ‘discipline, problem or task oriented’ grids, and the development of common services that can support their simultaneous operation and inter-operation.

The kinds of future functions envisaged are usually described as: on-demand, ubiquitous access to remote computers, digital repositories, scientific instruments and most recently sensor networks; virtual environments and on-line collaborative interaction; seamless resource sharing; cooperative working; and coordinated problem solving within dynamic virtual organisations.

The program will install, configure and operate servers and services that allow grid like workflows to use resources anywhere in the country, subject to suitable authorisation of course. This kind of capability is essential for research in disciplines such as high energy physics and climatology, but is also needed by NCRIS capabilities, such as AusScope, eMII and NANO.

As a general observation, at present, most disciplines are intent on developing standards and functionality that meet their own needs, under the pressure to get results, and have low tolerance for the kinds of timescale that would be needed to agree generic middleware or user interfaces. Australia therefore faces the problem of playing catch-up as global disciplines have sufficient funding to develop their own infrastructure components and effectively ‘go their own way’.

Australia could create its own unique and different national grid flavour or single out one of the international flavours of grids over the others and support and develop it. Neither option is particularly viable given the scale of investment required in the first case and the unsurmountable discrimination that would arise at a disciplinary level in the second case.

There is however some value in adopting and supporting only a limited set of grid flavours, especially as international efforts are working towards the inter-operation of

a number of major grid flavours. In particular, projects exist to enhance inter-operation between globus (US), glite (UK) and unicore (European) grid middleware.

The last point aligns with some very clear feedback from consultations in this area, which was that it would be best to focus on providing core services which are certain to be of value and to do them well. This is often seen as the only possible response to such a complex and broad agenda and the limited nature of funding.

APAC has developed a working grid infrastructure, and included proposals for grid services to its review. It included the full range of services commonly deployed in grid environments and included:

- **Authentication, Authorisation and Accounting:** A set of foundation services that identify users, authorise access to resources, manage the transfer of rights between users, and that report on and control resource use.
- **Information Services and Resource Discovery:** A layer of services that allow resource providers to publish their resource characteristics and availability and which allow users and workflow systems to find appropriate and accessible resources.
- **Scheduling and Job Management:** A set of system services that support the scheduling and execution of work on to available computing platforms.
- **Data Management and Movement:** A set of system services that support user data management across participating resources.
- **Application Support:** Components that support domain or application specific portals, workflow management, and value-add data services (such as metadata or content-based image location and retrieval).

However, the investment plan repositions most of these as follows:

- Wherever possible, authentication services will migrate to the Australian Access Federation, this program will only operate services which cannot migrate
- Information Services and Resource Discovery, and Applications Services will be pursued either as part of activities within other programs or within domain specific application activities mounted by other NCRIS activities. This program will provide the servers on which they can operate
- The program will implement a data movement service, data management services will only be installed as required by ANDS, or application and discipline services

From that, five key contributions can be identified at a cyber-infrastructure operational level, which application oriented workflows and user oriented tools and portals need, and which can be operated as services in common.

Therefore the proposed initial activities of the Interoperation and Collaboration Services program will be:

- A middleware operations function that can support several middleware and portal flavours operating over an integrated resource allocation and authorisation system
- A job submission fabric that allows the automatic unattended submission of jobs and workflows on any system within the fabric
- A data movement fabric that allows the automatic unattended transfer of data including effective large bulk transfers between any systems within the fabric
- A set of easy to use virtual organisation services, including AAF integrated collaboration workspaces
- The day-to-day operations management for the above services
- A small activity responding to requests for help and tracking requests through to expertise groups and services and their response

Investment into these services will include the provision of systems and servers, software and software environments, and the management means to ensure they are reliably configured and regularly tested for compliance and operational status.

The intention is that as additional requirements emerge from the other programs, and as responding services become clearer, they will migrate into this 'operational arm'. Therefore participants in this program might supply federation services and data retention capacity for ANDS.

Components developed in the Applications and Discipline Services would also migrate as they become robust and use reliance on them increased.

The Australian National Grid (ANG)

The ANG builds on APAC investments to date in grid systems, and extends those by providing for two classes of program participants, program members and affiliates. The goal is to have resource in at least one program member in every region.

Program Members will generally be service providers supporting compute and data services to researchers across multiple institutions. As program members, they will install and operate core ANG services on their premises and commit to configuring their own compute and data services to integrate with the ANG protocols and standards.

Program Affiliates will be organisations that seek to allow access to their research resources through ANG services or to seek access to other resources through ANG services, but will not be operators of the ANG services. Program affiliates will agree to configure their systems to comply with ANG protocols and standards and the regional program member will provide assistance as required. That assistance may extend to co-managing servers on program affiliate sites if that is required for reliable ANG operation.

The identified regions, with identified members and potential affiliates are:

Region	Member	Affiliates
ACT	ANU	GA, CSIRO, NLA, AA, ABS
NSW: Kensington	Ac3	Universities, ANSTO, 5.3
NSW: North Ryde	<td>	ATNF (5.10), AAF, MacquarieUni, CSIRO
QLD: South East	QCIF	Universities, IMB (5.1), ACCESS (5.13)
QLD: Far North	QCIF	JCU
South Australia	eRSA	Universities, 5.2, 5.3
Tasmania	TPAC	Uni Tasmania, CSIRO, eMMI
Victoria: Parkville	VPAC	Universities, Parkville Medical Institutes, BoM
Victoria: Clayton	<td>	Monash Uni, CSIRO, Synchrotron
Western Australia	iVEC	Universities

Significant work remains to be done to fill this table out to more properly represent the key targets for research resource integration with the ANG.

NRICC will:

- Establish an inaugural Steering Committee to provide guidance on the further development of the Australian National Grid (in consultation with DEST)
- Develop the program agreement and an initial program plan building on the existing APAC agreement and grid program plan (in consultation with DEST and the steering committee)
- Invite expressions of interest in the role of contract agency, program members and affiliates (in consultation with the Steering Committee)
- Finalise the program agreement and program plan (in consultation with the Steering Committee and participants)
- Appoint a Director: Australian National Grid, who will also be the Interoperation and Collaboration Services Program Manager within eRI.

The Australian National Grid will:

- Receive funds from NRICC under the arrangements defined in the program agreement and against the activities described in the initial program plan
- Operate as a single co-operative activity across all members and affiliates
- Engage the Steering Committee and others in developing a community of interest around grid technologies and interoperation standards and services, and engage in relevant international developments
- Develop and sustain the agreed services for middleware support, job submission, data movement, and VO services
- Establish and sustain an operation centre able to track and manage configuration requirements, identify grid performance, rectify faults and provide expertise to support infrastructure builders and operators in other programs
- Establish and sustain a help service which can direct user queries to the expertise most likely to be able to assist and track responses and resolutions
- Manage resources within the terms of the program agreement
- Report to eRI on the implementation of the program plan on a regular basis
- Propose a revised program plan on an annual basis

The ANG will be led by a Director appointed by eRI, and include a number of system administrators appointed at each member site.

Program members will provide these staff with the necessary site privileges to configure and operate the servers required by the activities below.

While the ANG will be operated on a best effort basis, for which users will need to accept risks related to availability at any point in time, appropriate service level goals will be included in the program agreement.

The initial set of services to be supported by the ANG will be defined in the initial program plan. In subsequent years, those service can be redefined and new services added through agreed changes to the plan.

Middleware Hosting

The ANG will identify a small set of inter-operable middleware protocols (expected to include globus and glite) and ensure all program members support access to their own systems via those protocols.

The ANG will also install a small set of inter-operable data access protocols (expected to include GridFTP, SRM and SRB) and ensure all Program members support access to their own resources via those protocols.

The ANG will identify core application support services such as gridsphere and arrange for suitable installations to be available on the grid infrastructure.

Overtime, the program plan will be updated to nominate additional protocols and services that will form part of the core inter-operation support.

Job Submission Fabric

The ANG will select standards and information models in accordance with standard practice for the supported middleware, that allows any user with grid credentials and appropriate authorisations, to submit a job for execution on any grid enabled resource at any member or affiliate site.

The ANG will be responsible for the standardisation of job submission semantics and protocols across the members and affiliates.

Data Movement Fabric

The ANG will select standards and information models in accordance with standard practice for the supported data access protocols, that allows any user with grid credentials and appropriate authorisation to access, replicate, and relocate any data associated with any member or affiliate resource.

The ANG will be responsible for the reliability and speed of data movement services.

User Collaboration Services

The ANG will provide a range of user collaboration tools including AAF integrated wiki's, AF integrated shared workspaces, and AAF integrated email lists.

The ANG will assist support access grid systems at each member site.

Operations

The ANG will establish an operations centre that can assure the reliable inter-operation of the participating systems and sites.

This activity will be responsible for the deployment and run-time management of servers, such as gateways and compute and data access servers. It will also provide a web based virtual 'operations centre' that provides information on the operational status of the infrastructure in real time and historically.

User support

The ANG will provide a "User support" function as a contact point, a web and email presence, and a ticket issuing and tracking system.

This system will be configured to lodge and track issues and problems as they arise and are dealt with across all programs under NRICC.

Reports from this service will be an input to performance assessment for all programs.

Budget

The program agreement will include:

- the arrangements and responsibilities for the program steering committee and the program manager.
- appropriate service level agreements
- an initial schedule assigning responsibilities to proposed participants and the process by which that schedule can be varied over time

Program Members will be regional service providers which support compute and data services to researchers relevant to NCRIS (and others) across multiple institutions. It is not a funding stream for internal IT services or for high end computing or data systems located within operational units.

At the financial level, NCRIS will support specific staff and systems functionality related to nationally co-ordinated core data and computation services. Service providers will supply additional resources to extend those to meet local needs.

A decision needs to be made on the eligible regions, noting that APAC allowed one per state and that perhaps this raises difficulty in the more populous and highly dispersed states (such as Queensland).

Such service providers will need to demonstrate their ability to support compute and data services, their ability to engage regional research institutions and communities, and undertake work to connect key research facilities in their region into the e-Research infrastructure.

An initial budget is under development.

Program: Applications and Discipline Services

This program area will bring together user, platform and discipline groups to work towards the vision:

Users are able to work together with familiar tools despite accessing multiple and remote compute, data, analysis, sensor or instrument resources.

This is an ambitious goal if translated across all possible research interests and research disciplines, so that this investment will focus primarily on NCRIS priorities.

However, the underlying motivation is that improved e-Research adoption depends on higher levels of content, application and system inter-operability and the provision of common tools and services that add value to research communities. Without tools and content inter-operability, researchers can only create bespoke components and isolated data sets that require costly case-by-case integration at a later date.

Because of the limited funding available, the program will focus on four broad goals:

- First to harden the definition, specification and implementation of the core middleware services required for data and compute system inter-operation
- Second to develop collaboration services that will support the broad e-Researcher community and then harden these for deployment within the national grid, an initial set of basic collaboration tools
- Third to work with ANDS to identify, implement and harden middleware and tools relevant to data curation and management for new e-Researchers
- Fourth to work with NCRIS capability and other discipline communities to identify the core middleware services they require, develop operational forms of these services and move them into the national grid

These are roughly in chronological order as the first is urgently required to stabilise extensions of the existing APAC grid infrastructure to a wider set of resources, and the last depends on further development of the NCRIS capability investments to clarify their cyber-infrastructure requirements.

This program is largely expertise based and includes a majority contribution of effort from participants, so that platforms for collaboration funding will be highly leveraged.

The leverage arises, because Platforms for Collaboration proposes to fully fund and manage the activities within the national grid and the regional organisations in which those services will be located, have agreed to contribute their effort to this program.

Similarly, where NCRIS may fund the on-going operation of CSIRO's re-configurable network, CSIRO will contribute effort into this program.

The rationale for this approach from NCRIS' point of view is straightforward. The service requirements from research communities are dominated by the evolving international discipline developments in middleware and tools, and only those intimately involved in such activities can offer advice on the directions likely to arise.

A gap exists and will continue to exist, where service providers are focussed around system and middleware issues, inherently seeking generic services, and research communities are pulled towards specific middleware solutions and highly specific user community tools by international developments.

The goal of this program is to close that gap by bringing together effort from service providers, from researchers in infrastructure services, and from researchers in application spaces. Without this program the continuous evolution and development of the operational cyber-infrastructure will lack essential guidance.

The rationale from regional service providers and research groups, is that separate piecemeal approaches to every NCRIS capability (as an example) is most likely to alienate everybody involved and harm their own goals.

Hence, while the regional service providers are funded to undertake such work and middleware oriented research groups are enhanced by engagement with applications, NCRIS can achieve significant overall benefits by developing a national orientation that engages these groups in coordinated (and preferably concerted) activities.

The particular importance to NCRIS is that the process is structured to identify, trial and test the middleware and user tools that can be adopted as part of the developing national cyber-infrastructure. The intention will be to migrate such tools into services provided by the more operational programs.

The challenge for the program is to manage variations as it undertakes an exploration of viable near term services working with program members who will all prefer committed funding and with application groups who will have independent priorities.

Unlike the other programs, the applications and discipline services program will not be implemented through a lead agency and a program agreement, as the bulk of the effort is and should remain under the management of 3rd parties.

However, the application of the \$12M funding to activities established in conjunction with sector participants will be approved by NRICC.

Three kinds of activities are contemplated, which may each involve different arrangements and timescales:

- Projects that ‘adopt, adapt and influence’ selected middleware components into demonstrable user and discipline focussed tools and services
- Investments in nationally significant expertise that adds value to infrastructure development and aligns with strategic needs identified by NRICC
- Resources for service providers to sustain additional tools and services as they are deployed into the operational infrastructure

The funding in these will vary over time, as will the relative emphasis between general user tools and more generic discipline services. Resource allocation will be based on the opportunities that are most likely to create new services or to enhance existing services within the operational infrastructure.

Establishment steps

NRICC will:

- Establish a National E-research Architecture Taskforce (NEAT) to provide guidance on the further development of the program of activities (in consultation with DEST)
- Appoint the program manager in Applications and Discipline Services, who will be the chair of NEAT (in consultation with DEST)
- Recommend members of NEAT some of whom may be appointed on a funded basis
- Approve activities in this program area as part of the NRICC business plan and from time to time during the start up process

The program will:

- Work through NEAT, with other NCRIS capabilities and identified discipline areas, the ANDS, ANCF and ANG, to develop a timeline for further infrastructure development

- Develop activities on a regular basis, noting that where activities are proposed over multiple years, annual reviews and justifications will be required
- Recommend projects, participants and funding arrangements to NRICC

Demonstrators

This is the primary means by which new tools and services will be identified, trialled and developed for deployment as an operational service.

Because of the steep learning curves in middleware, the number of technologies that need to be integrated to develop a service, and the need to develop understanding between communities, projects are expected to be generally of two years duration.

Projects will only be considered in response to the infrastructure development timeline published by NEAT.

Projects that deploy resources which do not require additional funding, can be approved by the executive director.

Projects which commit further funds, will need approval by NRICC.

Expertise development

Activities in this category will build nationally recognised centres of support for key technologies identified by NEAT. That expertise will assist adopt, adapt and influence global standards, components and middleware to create discipline oriented formats, tools and services, and would assist global peers in developing standards, components and middleware where appropriate.

Some of these centres are likely to be associated with the organisations involved in other parts of the investment plan, however some may be associated with e-Research ‘user’ communities. Proposals for expertise support will be approved by NRICC.

Tools environment

As new infrastructure components and particularly as an increasing variety of user tools are made operational, funding for additional expertise and operational cost needs to shift to the operators of the infrastructure.

Recommendations in this category will be negotiated with relevant services providers (most likely within the ANG) and would result in a change to the budget in that program to support newly developed services.

Foundation Services

The goals in this area involve the fundamental internet and human interconnects that make collaboration possible, through the two emergent capabilities:

An inclusive, high speed low latency toll free network, connects all researchers to all research resources, including international sites.

Researchers control who can do what with their resources through a common method based on a single login for each researcher.

The collaborative nature of e-Research and much of the activity that is envisaged under NCRIS will be significantly dependent on an ability to share data, systems and research infrastructure between research organisations and researchers, both within Australia and internationally.

At present, sharing data, systems and research infrastructure, in a secure way, between various universities and including research agencies presents a very difficult problem, both technically and from a policy perspective.

Therefore the two challenges contemplated in this area are as follows.

- To confirm the concept of the AREN and to implement steps to align the backbone, regional and campus network system elements with that vision
- To build a national access control system that supports effective collaboration between institutions and secures resources and access to known users

These services, have a much broader use than research, and while discussed here in terms of research objectives, will add value across the research and higher education sectors for many other uses. The proposition is that they be funded by the institutions gaining that value, to better direct their wider development.

The Australian Access Federation

Recent technical advances provide a new solution to secure sharing of data and systems using a “Trust Federation”. A Trust Federation is a combination of policies and technologies that allows individuals across many different organizations to securely share and analyse data within a trusted environment.

Some of the very significant advantages of a Trust Federation are that a researcher needs only one name and password (or its equivalent, such as a digital certificate) to access all appropriate resources at different institutions and that these credentials

would be issued by the researcher’s home institution. Additionally, an institution will need only one agreement with the trust federation rather than 40 or 50 individual agreements. The federation can also help to reduce costs through the collaborative purchase and/or support of some tools and technologies.

Previous DEST funding on Trust Federations through the Systemic Infrastructure Initiative (SII), particularly the MAMS (Meta Access Management System project (based at Macquarie University) and the first phase of the e-Security Framework project (based at the University of Queensland) have provided a foundation for an Australian Trust Federation. For instance, the MAMS test bed Trust Federation already has 9 universities with over 600,000 members. The e-Security framework project has established a PKI infrastructure with interoperation between different universities, has obtained “in principle” agreement from the US Higher Education and Federal Government PKIs for interconnection with the Australian federation, and has made significant headway with Microsoft and other vendors in updating their browsers to support the Australian PKI Infrastructure.

The Australian Access Federation will be established as a stand-alone operational entity as it will need to enter into a wide range of commercial contracts and as its services need to be developed and scoped against a very wide potential market.

.Nevertheless, the AAF is being established on the back of commonwealth funding and NCRIS will provide continuing support during its development phase.

The AAF will provide PKI and shibboleth services to the entire research community and host a National Centre of e-Research Expertise in those areas.

The AREN

Support is provided in this plan to sustain the link to Darwin beyond (2008).

The program activities should also assist the AREN by bringing more researchers into its scope, and the ANG and ANDS programs can be expected to significantly increase research on-net traffic, assisting the value proposition of the AREN to institutions.

A further issue that will be addressed is the provision of a source of design and implementation advice that can harmonise campus infrastructures with the AREN. An activity under the expertise component of the application services to develop and supply that expertise is envisaged.