

BeSTGRID Middleware to support New Zealand eResearch

Executive Summary

As research infrastructure in New Zealand begins to utilise advanced computing platforms and networking, a remaining challenge is to make the benefits of such resources accessible to different research communities that might be physically dispersed throughout the country—bringing the power of advanced computing to the research desktop. To do this requires a significant investment in software, to tackle problems as diverse as authentication and access rights through to running complex simulations across a range of cluster computers. Grid Middleware is the set of software services and protocols that integrates individual computers, storage systems and application tools into a coherent, distributed, computational environment. Such environments are critical in allowing NZ researchers to scale up to larger or more complex problems, to share their resources and results more effectively and to connect with established offshore eResearch communities.

This proposal lays the groundwork to allow researchers in universities and CRIs to effectively incorporate shared computational resources and data into their research activities. We describe a plan for the ongoing support of BeSTGRID (originally funded by TEC) as the organisation to carry this work forward. We propose establishing a solid middleware infrastructure that is: (i) carefully selected from the technologies available, (ii) well managed, (iii) made sustainable via replication and education, and (iv) demonstrated via engagement with key research groups in the geosciences and biosciences. The emphasis will be on delivering services that are generally useful to many communities, but with the specific examples used to test end-to-end functionality in two research areas that are critical to New Zealand's economic and environmental well-being (geosciences: *SCENZ-Grid* at Landcare and GNS, *AuScope* in Australia, biosciences: the *Virtual Institute of Statistical Genetics* and *NZ Genomics Ltd* both pan-University/CRI organisations). These groups will be used as exemplars to test and evaluate our middleware and to raise awareness of the capabilities and benefits of BeSTGRID as a strategic national e-Research initiative.

The programme of middleware development work is outlined below in Tables 2 and 3. Some of the specific technical tasks will be undertaken in partnership with the Australian e-Research programme to foster better collaboration, and to place New Zealand in a position to capitalise on the massive current investment in eResearch across the Tasman. To ensure sustainability we will also develop training resources and opportunities, and engage in outreach to encourage wider uptake beyond the explicitly supported projects and communities. While the majority of the budget requested is for software deployment and middleware service provision, our plans also emphasise the need for careful technology scanning, community sustainability and good governance as follows:

1. Coordinate national strategy for grid middleware development, making use of best-practice grid developments offshore and building on what we have already in place in NZ.
2. Sustain the established grid middleware community and provide it with the resources to be successful in supporting NZ research.
3. Provide a level of functionality above the command-line to make national grid computing platforms accessible and easy to use for distributed research groups.
4. Migrate selected research domains to the grid computing infrastructure.

Steering and technical committees are already established and operational for BeSTGRID; what it lacks is the financial resource to engage in coordinated strategic service development and provision. The lead investigators will work with the steering and technical committee to select the best technologies and international collaborations that can be used to develop our grid infrastructure. The project team will be drawn from researchers and IT support staff at universities and CRIs throughout the country.



BeSTGRID Middleware to support New Zealand eResearch

Principal Investigators: Mark Gahegan, Nick Jones, John Hine, Michael Black

A. Profile of BeSTGRID

BeSTGRID has established initial grid infrastructure for research and science in New Zealand, via support from the TEC IDF fund from late 2006 through to early 2008. The TEC contract was led by Auckland, with subcontracts held with Massey and Canterbury universities. BeSTGRID now coordinates grid activity across a partner network including University of Auckland, Massey University, University of Canterbury, University of Otago, Victoria University Wellington, GNS Science and Landcare Research. These institutions are currently (or planning to become) resource providers to BeSTGRID and most are active members of BeSTGRID's governance. Through support of member institutions' Office's of Research Management, IT Directorates (and in some cases dedicated Centres of eResearch), BeSTGRID continues to run as an informal collaboration with participation from current and new member institutions. Many members reinforce their commitment by openly and freely sharing their computing cluster and data storage resources.

This proposal will be led by an experienced Science and Technology team contracted from the University of Auckland, comprised of Prof. Mark Gahegan, Mr. Nick Jones, and Mr. Richard Li. Subcontracted institutions include the BeSTGRID partner network institutions, along with additional institutions able to contribute in specific areas. At each institution there exists senior academic and administrative support for the proposal, and experienced and skilled technical staff seconded to carry out the programme of work. Further specifics are covered under the section 'E'. titled 'Fund Management'. Biosketches describing the relevant experience of the principal investigators appear later on [Page 9](#).

The BeSTGRID community does not see its membership as being limited to its founders, and our proposal includes specific goals to: (i) broaden participation, (ii) include other university and CRI members and (iii) add in additional research communities and their tools.

B. Grid experience

BeSTGRID aims to establish and supply the infrastructure required for a shared distributed resource model for research support both in New Zealand and to support international research collaborations. Achieving this aim requires establishing an infrastructure of loosely coupled technologies described by grid computing and implemented in support of eResearch. The services offered within BeSTGRID so far have been focused on top-down science drivers (Bioinformatics, primarily Phylogenetics), and bottom-up service development and capability building.

Due to the immaturity of eResearch activity in New Zealand, BeSTGRID's approach over the last 2+ years has been to identify best-of-breed technologies developed internationally which are strongly aligned with local research and science needs, to implement these pragmatically, and to foster strong collaborations with researcher user-communities and eResearch technology communities locally and internationally. This has required much outreach and engagement, which has led to broad awareness and a scope of activity beyond the preliminary Bioinformatics focus. This, in turn, has provided substantial insights into the needs for eResearch infrastructure in New Zealand.

BeSTGRID has established the basis for a shared eResearch infrastructure through implementing a variety of middleware. With a focus on acquiring mature configurations of technology from the international community, BeSTGRID adopted architectural principles from the Australian eResearch community, based on developments in Ireland, the USA and elsewhere. The architecture incorporates components in support of science application portals or gateways, identity and access management, job submission and management tools, storage and computing resource control, operational monitoring and management of resources, and frameworks for development and extension of core technologies for integration and subsumption. Specific technologies with which



the BeSTGRID Technical Working Group has already gained experience include (but are not limited to):

- Job Submission & Control (PBS, WS-GRAM),
- Program Compilers (GCC, IBM's C compiler),
- Identity Management (Shibboleth, x.509, Short Lived Credential Service, MyProxy),
- Virtual Organisation Management (low level) (GUMS, VOMRS),
- Globus Toolkit (to help construct basic grid middleware services),
- Grid Data Services (Storage Resource Broker, GridFTP),
- Computing Cluster Management (Rocks, Condor),
- Virtual Community Technologies and Applications (JIRA for software issue tracking, SAKAI for community communication, blogging and sharing of ideas and plans) and
- Development Frameworks (CoG Kit, GridSAM, GridSphere Portal framework, Grisu command line).
- Operations Management systems (GoC, INCA, Nagios, WebMDS, MRTG)

As an established and working community, BeSTGRID is actively engaged in the task of supporting this nascent eResearch infrastructure and providing national leadership in establishing the necessary middleware to support sharing of resources and research collaborations.

C. Understanding of requirements

“There are four areas – distinctive to New Zealand – where new knowledge and effective innovation are vital for our success. In these areas, New Zealand RS&T must shine.”¹

The four areas identified by government as being vital for New Zealand's success are: (i) Drawing sustainable value from our biological base, (ii) Recognising and managing our environment as an asset, (iii) Seeding and supporting high-tech industries and (iv) Strengthening services for health and social wellbeing. Each of these areas spans many research disciplines and related knowledge domains, and increasingly requires sophisticated information management and analysis and complex modelling to innovate and exploit. The infrastructure of the modern researcher has evolved to be highly dependent on Information Technology, with the last 15+ years seeing the rise of technologies to support collaborative and distributed research, now loosely described by the term eResearch, and underpinned by advanced research networks and grid computing.

As research infrastructure in New Zealand begins to make use of high speed networking and advanced computing platforms, a remaining challenge is to make the benefits of such resources accessible to different research communities that might be physically dispersed throughout New Zealand, and the rest of the world—bringing the power of advanced computing to the desktop. To do this requires a significant investment in software, to tackle problems as diverse as authentication and access rights through to running complex simulations across a range of cluster computers. Grid Middleware is the set of software services and protocols that integrates the individual computers, storage systems and application tools of research communities into a coherent, distributed, computational environment. Such environments are crucial in allowing NZ researchers to scale up to larger or more complex problems, to connect with established offshore eResearch communities and to share more effectively their individual resources and results.

An effective, coordinated, strategic approach to national grid middleware infrastructure and service provision is needed for the following reasons:

1. Funding levels are not sufficient to take a fragmentary or piecemeal approach to providing the services that NZ research communities need. A national, coordinated approach to

¹ identified in the recent publication “Government's Agenda for New Zealand Research, Science and Technology”.



design, planning and deployment must be taken to ensure that all the pieces we build will fit together into a cohesive whole.

2. The resulting middleware needs to meet the structure and standards being developed overseas, so that we can integrate seamlessly with offshore communities. Grid middleware is a fast-evolving field, and most progress is occurring outside of New Zealand. It is therefore necessary not only to develop a national strategy, but also to work with international leaders in this realm to refine and update our strategy as new technologies and opportunities arise.
3. Those developing middleware need to be tasked appropriately, managed well and be held accountable for delivering the services they undertake to provide, otherwise researchers will not get the services they need.
4. The tools and services built need to address the real needs of research communities who are ready to work with them.

Leveraging our base of experience, BeSTGRID seeks to mature current services and extend the community of participants and contributors, as described in the detailed plans in relevant sections below. While there is a recognition of the many potential uses of shared eResearch infrastructure across the New Zealand science system, BeSTGRID's strategy makes careful progress towards a stable model for service provision across a shared and distributed resource base; a wider skillbase supporting this provision to sustain current capabilities; and lays the groundwork for the future. To ensure that the infrastructure being developed is shareable and generalisable to other research domains, BeSTGRID will support research-driven uses of the infrastructure, which requires implementation of shared research applications that make use of our middleware services.

C.1 Overall proposal strategy

The landscape of possible middleware development to support New Zealand researchers is vast; this proposal cannot possibly address it all. We've therefore adopted an approach that ensures a focus on building generalised infrastructure that is shown to be effective through community-derived use cases.

Figure 1 shows that the landscape can be understood in terms of layers of services, from the network, computational, and storage infrastructure at the lowest level, through to sophisticated research communities and their discipline specific applications at the highest level. The layers of middleware can also be characterised vertically according to the level of abstraction they provide from the complexities of the raw computing infrastructure. Between these two axes, a matrix is definable. This is described in more detail as relates to this project in a later section, in Table 2. We have carefully chosen specific functionality to deliver throughout this matrix. Our choices are based on: (i) sustaining and building on current capabilities (iii) the drivers and needs of specific research communities, and (ii) remaining compatible with best practice overseas, and most especially in Australia.

In sustaining and building on current capabilities, the middleware to be implemented will have strong alignment with existing community needs and current technology choices. Implementation and support will be carried out by the existing BeSTGRID developers, augmented by new members of the community from current and new sites committed to BeSTGRID and this proposal. To ensure that the middleware truly meets the needs of research applications, and is driven by them, we plan to develop a small number of tools and services at the upper levels of the matrix in Figure 2, in specific research fields as shown.

In selecting the specific research communities to be supported, attention was paid to strategic and transformative areas of research, as identified in the recent MoRST publication "*Government's Agenda for New Zealand Research, Science and Technology*", and with compatible activities supported by the REANNZ administered KAREN Capability Build Fund. The specific programmes selected to demonstrate the effective implementation of grid middleware include SCENZ-Grid (KAREN CBF), the Virtual Institute of Statistics Genetics (TRST, NERF), and New Zealand



Genomics Ltd (RIAG). The sparsity of research programmes selected indicates a potential to identify a much broader set of research capabilities that could be supported in future.

We note that a far more comprehensive funding program is urgently needed to address the full extent of eResearch needs of these and many other communities; to provide them with access to the resources they need; and to work directly with these communities to migrate their work practices into the emerging eResearch infrastructure.

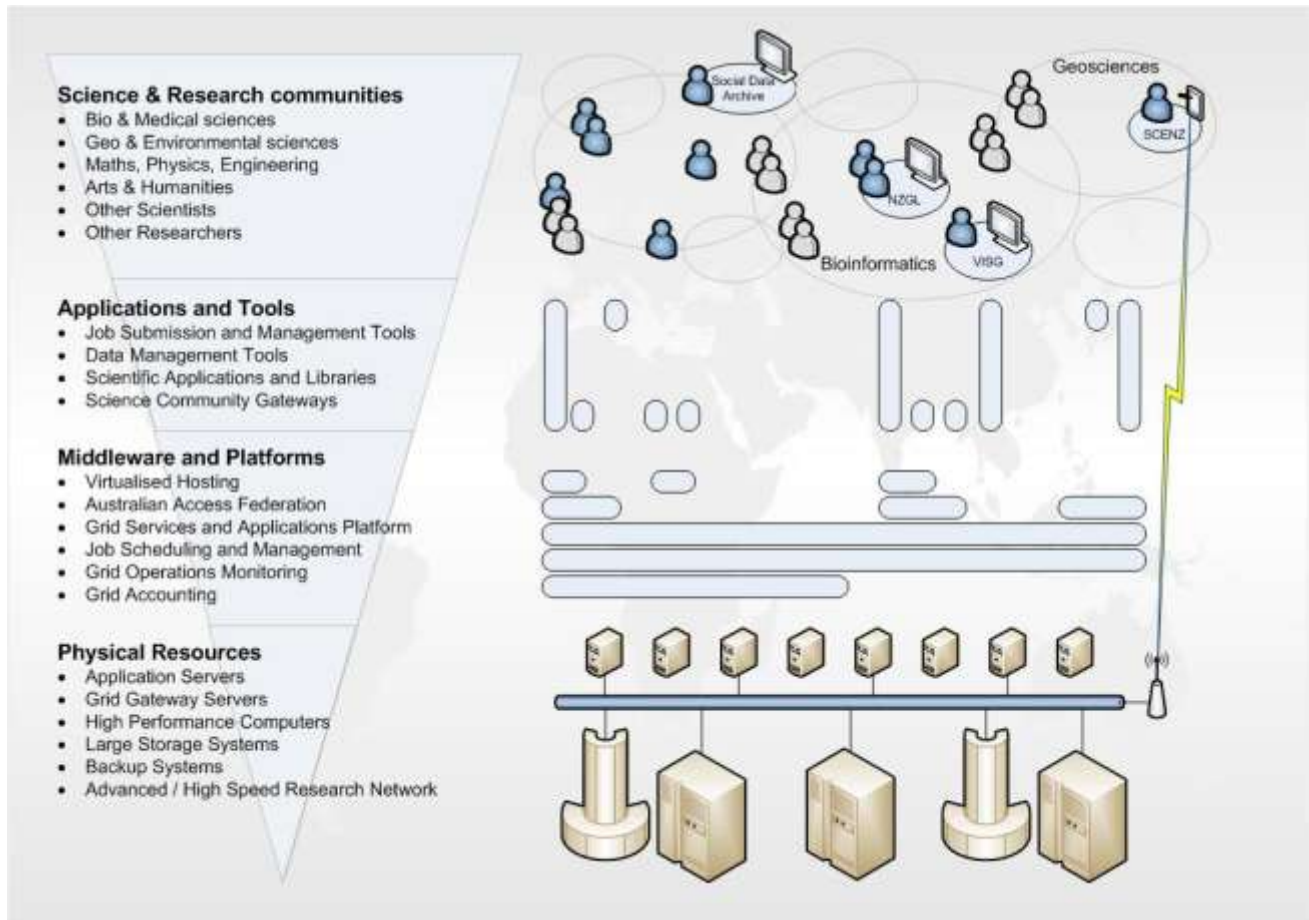


Figure 1: Schematic of eResearch landscape approached within this proposal

D. Funding recommendations

D.1 Description and rationale

This proposal recommends building on the considerable achievements of the BeSTGRID community, to further evolve and sustain the BeSTGRID grid infrastructure in support of the New Zealand research sector. The proposal describes funding in 5 specific areas:

- (i) **Middleware deployment, integration and testing,**
- (ii) **Grid middleware strategy development,**
- (iii) **Community sustainability,**
- (iv) **Program management and**
- (v) **Community leadership and governance.**

A summary of the activities involved in each of these major areas is given below in Table 1. This is followed by a detailed description of each area's proposed activities.



Table 1: major project areas and details of associated activities.

Area	Description
Middleware deployment, integration, testing & development	Includes deployment and customisation activities—which will form the majority of the work with development of entirely new infrastructure minimised or avoided where possible. Coordinated operations and service management.
National grid middleware strategy development (scanning, planning, acquisition, collaboration)	Technology and community scanning, roadmapping, compatibility and functionality review, selection, operational planning. Outreach to offshore eResearch communities, especially the Australian ARCS initiative, the e-Science community in the UK and some of the cyberinfrastructure projects in the USA.
Community sustainability	Replication of expertise and functionality at different sites, training & outreach activities, developer materials, workshops, opportunities for obtaining additional partners / research communities.
Project management	Includes managing sub-contracts for development and day to day project management, selection of project management tools and strategies, ensuring quality and timeliness, reporting on progress to MoRST and wider community.
Community leadership and governance	Grid middleware leadership, research community liaison, high level coordination with KAREN/REANZ, MoRST, Universities and CRIs.

MIDDLEWARE DEPLOYMENT, INTEGRATION, TESTING & DEVELOPMENT

Middleware deployment will be the largest budgeted item and details of the tools and technologies we plan to deliver are given later. This work will be undertaken by individuals and groups at various BeSTGRID institutions via sub-contracts. We specifically plan to avoid developing new software tools unless absolutely necessary, relying instead on adopting best practice tools and services from offshore where possible (see Strategy Development). This will save time and money as well as minimising the risks associated with sustaining bespoke software in a small community,

Table 2 below describes an overview of the infrastructure areas being targeted, from a research domain and technology perspectives. Our workplan starts from the infrastructure already developed as part of BeSTGRID, and shown in the second column of the table. We then extend this into generic tools and services, and into infrastructure areas of critical strategic importance to the New Zealand research landscape, leveraging BeSTGRID members’ science disciplines. This provides a strategic direction while managing risk by ensuring strong engagement with science leaders.

Table 2: Middleware infrastructure investments by science and technology

Infrastructure areas	Existing Infrastructure	Generic researcher Infrastructure	Strategic Bio-sciences	Strategic Geo-sciences
Value drivers	Sustain Capability and Maintain Quality of Service	Broad support for transformative capabilities Generalised middleware and tools for broad use	Strategic science community support (NZGL – RIAG, VISG – TRST NERF) Support to transform community	Support for capability investment in spatial geosciences (SCENZ-Grid) and in seismic information services Support to transform community
Skills development	Technical Working Group	Workshop programme Grid summer school Conference attendance	Workshops or Summer School with domain conference/group	Workshops or Summer School with domain conference/group
Workflows	SCENZ-Grid	Exemplars and templates Taverna knowledge	Taverna workflows published in MyExperiment	Kepler workflows Kepler SRC actors Taverna spatial



BeSTGRID Middleware to support New Zealand eResearch

		Kepler knowledge MyExperiment knowledge		workflows published to mySCENZ-Grid (based on MyExperiment)
Application Services	Biportal, focused on Phylogenetics	Survey generic user needs, resulting in e.g. Matlab and R libraries and exemplars	Bioinformatics libraries Biochemistry libraries	Geospatial libraries GEON portal OGC WFS, WPS, WPS-G mySCENZ-Grid collaboration portal
Virtual Organisations	Project team or researcher specific virtual organisations primarily in Bioinformatics	Guidance on drivers and how to use Setup per research group Use to control access to resources	Bio related VOs	Geo related VOs
Job Management	GRISU Command line tools CoG Kit Biportal	Exemplars and templates GRISU exemplars CoG Kit exemplars Globus WSRF exemplars	Biportal exemplars	GEON exemplars
Storage and Data services	GridFTP 2 sites SRB 1 site	Guidance on drivers and how to use GridFTP for large replicas Evaluation of SRB vs iRODS	SRB/iRODS Bio Catalogue	SRB/iRODS Geo Catalogue
Hardware Resources	4 compute clusters 3 storage sites 4 grid gateway sites 2 virtualised application farms	Guidance on how to access and how to contribute to BeSTGRID Review of resource allocation schemes, accounting and charging models	Auckland Bioengineering Institute cluster NZGL Cluster(s) and storage	Landcare Research cluster and storage
Operations Management	ARCS GoC INCA WebMDS Nagios	Test driven service monitoring with INCA Coordination of service issue management Support & feedback channels		

Whilst the grid computing model describes a shared and generalised infrastructure, there may be unique configurations required for each research domain. Specific projects come from the intersections of each of the research domains with the infrastructure technologies in the table above. Within each science domain there are many dependencies vertically through the infrastructure layers, with projects defined to gain experience across all layers. This approach ensures that end-to-end capability is created, laying the groundwork for future integrated approaches to infrastructure development.

Table 3 details the suggested technical programme of work that we plan to undertake, and the people involved. But since grid middleware is a fast-changing field, ongoing technical scanning and planning will be used to identify and integrate the best available tools at all levels (see *National Grid Middleware Strategy Development* below).



Table 3: Grid Middleware development personnel and tasks / responsibilities

Person	Tasks / Responsibilities
Nick Jones (Auckland)	Overall technical leadership
Vladimir Mencl (Canterbury)	OGRE, SRB/iRODS, Porting genetics software to grid, bio-portal, Scientific Application Libraries, sustainability
Kevin Buckley (Victoria)	Scientific Application Libraries, Virtual organisations, certificates, Geoscience Workflows, Grid gateways, GridFTP, science applications, sustainability
Paul Grimwood (GNS) & Robert Gibb (Landcare)	Geoscience Workflows, SCENZ-Grid, sustainability
Yuriy Halytsky (Auckland)	SRB/iRODS, GridFTP, GRISU, Bioscience Workflows, Shibboleth, Semantic mediation service, Scientific Application Libraries, sustainability
Niels Hoffman (Landcare)	SCENZ-Grid, workflows, GIS grid processing web services, sustainability
Chris McDowall (Landcare)	Client-side geospatial visualisation (in kind)
Kevin Richards (Landcare)	Biodiversity semantic mediation web architecture (in kind)
Andrey Kharuk (Auckland)	Scientific Application Libraries, Virtual organisations, certificates, sustainability
Stewart Stevens (Otago)	Scientific Application Libraries, Virtual organisations, certificates, sustainability
IT Developer / sys-admin (Massey)	SRB/iRODS, GridFTP, GRISU, Bioscience Workflows, Shibboleth, Semantic mediation service, Scientific Application Libraries, sustainability
Andrew McBeath (NZSC)	Gateway development, Virtual organisations, certificates, GridFTP, sustainability

NATIONAL GRID MIDDLEWARE STRATEGY DEVELOPMENT

Our technical workplans, as outlined above, represent our best understanding of the landscape of Middleware technologies at the time of writing. But this is a very dynamic field, where new developments are announced weekly and opportunities for shared development are common. A national strategy for middleware development must be responsive to (a) the current situation (b) future technologies as they arise and (iii) opportunities to share or borrow development effort where it will save valuable time or other resources.

The main work items in this area are technology scanning, community needs assessment, technology roadmapping, compatibility and functionality review, middleware selection, and operational planning. Travel funds will be used here to investigate promising software and related technologies, work with NZ research groups as stakeholders and make and maintain a working relationship with other large scale, offshore eResearch/eScience/cyberinfrastructure projects.

At this level, coordination needs already exist in regard to several ongoing projects (such as identity management, storage schemes, and advanced video-conferencing) especially with the Australian ARCS initiative, the e-Science community in the UK and some of the cyberinfrastructure projects and government science laboratories in the USA.

This group will meet weekly, using video conferencing, to share their plans, knowledge and findings. As part of our middleware strategy development, we will also host a national workshop (with open invitation to all interested parties), possibly in conjunction with a local High Performance Computing or similar conference. We will time this towards the middle of the project, after some initial scanning and investigation has been completed. Table 4 shows project personnel will be included in middleware strategy development.



Table 4: Personnel involved in middleware strategy development

Person	Tasks / Responsibilities
Nick Jones (Auckland)	General strategy, liaison with the Bioscience community & ARCS, Argonne National Lab, Monash University
Mark Gahegan (Auckland)	General strategy, liaison with the Geoscience community & ARCS San Diego Supercomputer Center, eScience Centre at Edinburgh and Manchester
John Hine (Victoria)	General strategy, liaison with the Geoscience community
Mik Black (Otago)	General strategy, liaison with the Bioscience community
Paul Grimwood (GNS)	Coordination with GNS / SCENZ-Grid
Yuriy Halytskyy (Auckland)	Underlying Grid technologies
Robert Gibb (Landcare)	Coordination with Landcare / SCENZ-Grid
Vladimir Mencl (Canterbury)	Underlying Grid technologies
Stuart Charters (Lincoln)	General strategy, liaison to other (non-BeSTGRID) communities
Tim Molteno (Otago)	Resource allocation schemes, accounting and charging models
Chris Messom (Massey)	General strategy, liaison to other (non-BeSTGRID) communities

COMMUNITY SUSTAINABILITY

Our sustainability strategy has two strands, (i) sustaining the middleware and development community within the project and (ii) sustaining and educating the national community more broadly.

(i) We see middleware sustainability as going hand in hand with the software activities described above, to ensure that national grid middleware is not dependent on a single site or individual. To meet this aim, to provide training opportunities for other collaborators and to help widen the skillbase in New Zealand, we will replicate each implementation we develop at least one other institution, with accompanying transfer of knowledge and skills. Each specific work item, when completed will be 'handed over' to other partner institutions. So, each developer in this project will also be devoting some of their time to community sustainability. The technical work table (Table 3) shown above lists this work item for all developers. We will achieve these aims via videoconferencing and use of community support tools such as Sakai for minor projects, and by face-to-face installation, training and workshops in other cases. Travel funds have been allocated specifically for this purpose (see budget).

(ii) Where needed, sustainability may also include knowledge and skill transfer beyond initial partner institutions. (Software developed will be kept in the open source domain wherever possible, be professionally managed and will be made available to other parties as needed.) Additional research communities that wish to engage in either service provision or middleware deployment and use will be supported as resources allow. For the more general community, we will deliver an educational workshop specifically targeted at junior researchers and developers, and organised jointly with the Australian ARCS group. This will be co-located with a suitable national conference. To increase national awareness and take-up, the Grid Middleware Strategy team (see above) will also engage in outreach and education as detailed in Table 4.

PROJECT MANAGEMENT

Project management includes managing sub-contracts for development and day to day project management, selection of project management tools and strategies, ensuring quality and timeliness, and reporting on progress to MoRST and wider community. General day-to-day activities will be led by an experienced project manager, Richard Li. Nick Jones and Mark Gahegan have considerable experience in managing large, distributed development teams (see Table 5).



Table 5 project management team

Person	Tasks / Responsibilities
Mark Gahegan (Auckland)	Contracts; Liaison with Research office, Universities and CRIs development team and management team
Richard Li (Auckland)	Technical program project management, operational administration
Nick Jones (Auckland)	Liaison with development team and management team

COMMUNITY LEADERSHIP AND GOVERNANCE

High level strategy and governance issues, and strategic decision-making, will be the responsibility of the project leaders (Table 6). Their tasks will include keeping the project on schedule, being responsive to REANZ, MoRST and other branches of government, and liaison with senior management at the universities and the CRIs.

Table 6. Leadership and governance team

Person	Tasks / Responsibilities
Mark Gahegan (Auckland)	Liaison with MoRST, NZ universities, government, geoscience lead
Nick Jones (Auckland)	Liaison with CRIs, government
John Hine (Victoria)	Liaison with REANZ (John is Chair of the KAREN capability fund)
Mik Black (Otago)	Liaison with offshore and onshore bioinformatics communities

Professor **Mark Gahegan** is Director of the Centre for e-Research and Professor in the School of Geography, Geology and Environmental Science at the University of Auckland. He recently relocated to New Zealand, following nine years as Associate Director of the GeoVISTA Center for visualization and spatial analysis research at Penn State (www.geovista.psu.edu). His degrees are in the computing sciences. His research interests are at the interface of the natural and computational sciences, concentrating on eResearch, knowledge representation, visualization and geographical information science. He has published over 100 articles, has been a lead investigator on 25 funded research projects to date, and is on the editorial board of seven international journals. Prof. Gahegan leads the GeoVISTA *Studio* (exploration and discovery) and *ConceptVista* (concept mapping and ontology) open-source software initiatives. He has been deeply involved in several cyber-infrastructure or eScience projects in the USA and Europe, including:

- the Geosciences Network (GEON: www.geogrid.org). a large-scale US eScience project for the geoscience community,
- Human Environment Regional Observatories (HERO: www.hero.psu.edu). a large-scale US eScience project for the climate change impacts community,
- Fungal Plant Pathogen Database (FPPD: <http://fppd.cbio.psu.edu/index.html>), phylogenetics of plant pathogens, linked to their geography and ecology and
- Learning Activities in Digital Libraries (DialogPLUS: www.dialogplus.org), infrastructure for sharing teaching and learning resources between the USA and the UK.

He has worked on national science strategy, as a member of two National Academies panels in the USA, and is currently a member of the MoRST *FutureWatch* Science Scanning Network (<http://www.morst.govt.nz/current-work/futurewatch/>).

Mr. **Nick Jones** is Director of BeSTGRID and Manager of the Centre for eResearch, and consults to industry on software engineering and computer science related research problems through the Centre for Software Innovation, at The University of Auckland. Nick has been involved in education, industry, and research related technology and capability development programmes for the past 8 years, after postgraduate study of IT innovations using anthropological research methods. Nick's



expertise is in formulating strategy around high tech IT (in particular based on Internet technologies), with hands on experience in managing research, development, and commercialization programmes. Nick joined BeSTGRID from his role as CTO in the University of Auckland's first IT based commercialization, e-Learnings, where he was one of two senior managers tasked with developing, marketing, and selling eLearning platforms into Health and Education sectors in New Zealand, Malaysia, and Thailand. Nick has been involved in BeSTGRID since 2006, and was made a KAREN Champion in 2007 in recognition of his contributions to the community.

- led 2 TechNZ TIF Expert Fellowship proposals and received substantial FRST PSAF funding to develop eLearning platform technologies for commercialization from within the University of Auckland;
- Chair of the BeSTGRID Steering Committee, and member of the University of Auckland eResearch Executive Committee; and
- Research Strategist with the Network Design and Simulation Group, based in Engineering Science at the University of Auckland. Involvement in research strategy and commercial strategy formulation.

Professor **John Hine** is currently Dean of Engineering and Head of the School of Mathematics, Statistics and Computer Science. From 1 January he will become Head of the new School of Engineering and Computer Science. John's research is in distributed systems, particularly as it applies to e-Science and e-Research. John has been active in the development of research and education networking in New Zealand since 1983. He played a significant role in the introduction and development of the Internet in New Zealand, for which he was made a Fellow of the Internet Society of New Zealand. He was a member of the steering group that authored *Collaborating at Speed: Innovation Infrastructure for a Knowledge Economy*, a report that led to the development of KAREN. John is currently Victoria's e-Research champion and Chair of MoRST's Capability Advisory Panel for KAREN. In 2007 John was made a KAREN Fellow.

- Director of NetLink, Wellington from 1997 through 1999 when it was sold to Telstra-Saturn. NetLink was a startup ISP arising from the Internet development at Victoria University in the previous decade. Director of DOMAINZ, Wellington from 1997 through 1999. DOMAINZ was set up by the Internet Society of New Zealand to manage the domain service for the .nz domain.
- Member of the Council of the Internet Society of New Zealand from 1995-1998 and 2000-2002. Member, Information Technology Advisory Group to the Minister of Information Technology, 1997-1998. Board Member, Next Generation Internet New Zealand from 2004 through 2006.
- In 2000 John chaired the Internet Society of New Zealand working group that replaced DOMAINZ with a "Simple Registration Service" that provided minimal domain name registration needed by all domain name holders while allowing a market of value added services to develop.
- Member, MoRST's ICT Infrastructure Advisory Group, 2003-2004. Chair of the MoRST sponsored Capability Development Advisory Panel for KAREN since 2006.

Dr. Mik Black is a Senior Lecturer in the Department of Biochemistry at the University of Otago. Dr. Black's primary expertise lies in the development and application of statistical tools for the analysis of data from high throughput genomics experiments. The nature of this work, however, which is both collaborative and computationally intensive, has led him to become involved in the provision of eResearch infrastructure for New Zealand's genetics and genomics research communities. These involvements include:

- a leadership role in a successful proposal to establish a national genomics infrastructure through MoRST's RIAG process, New Zealand Genomics Ltd (NZGL);



- co-PI on a successful KAREN CBF proposal, *Integrated Genomics Resources for Health and Disease* (joint with Dr Chris Brown);
- involvement in a successful bid to establish the FRST-funded *Virtual Institute for Statistical Genetics (VISG)*; and
- membership of both the BeSTGRID Steering Committee, and the University of Otago eResearch Advisory Group.

While seemingly disparate, a common theme across these areas is the provision of tools to allow New Zealand researchers to collaborate, share data and perform advanced data analyses. Although the projects listed under (1) and (3) above are still in their early stages, the *Integrated Genomics* project (2) is well-advanced, and has put tools in place for the KAREN-based analysis of both microarray and DNA sequence data by New Zealand scientists. These tools are currently being well utilized by researchers, as well as being used in bioinformatics courses taught at the University of Otago.

D.2 Amount of Middleware Fund required and any other funding available

We request the full amount available in the RFP, \$840,000.

We will utilise these funds over a fifteen (15) month period. In addition, we will leverage the following in-kind contributions or support, from our host institutions (Table 7):

Table 7: Complementary (matching) funding, contributions and support

Institution	Contribution or supporting item
University of Auckland	ITS services & support Geoscience PhD Researcher (Brandon Whitehead) KAREN Champion Auckland BeSTGRID Cluster
Victoria University, Wellington	Up to 0.5 FTE programmer position
Otago University	0.1 FTE senior researcher
Massey University	Massey BeSTGRID Cluster
Canterbury University	GridGateway and underlying virtual machines, DataGrid with its 20+ terabytes of capacity and access to the BeSTGRID share of the HPC and EEE C computer clusters
Lincoln University	KAREN Champion
Landcare Research	Semantic-Web and Geospatial Visualisation projects MyExperiment setup
GNS Science	KAREN Champion

Our rationale here is to keep funding levels high enough to ensure there is continual progress, and to make rapid progress towards supporting NZ researchers. However, we will begin the process with careful technology scanning and implementation planning, so the majority of the development work will begin after this has occurred (see Milestones for more details). We anticipate a smooth project initiation and early planning sessions as the project staff are already in place at their various institutions.

D.2.1 Budget details

The programme of work has the following major expenditures areas:



Table 8: Programme Areas, Leads, and Cost

Task	Lead	Cost
Middleware deployment, integration, testing & development	Nick Jones	\$410,000
Community sustainability	Mark Gahegan	\$100,000
Project management	Mark Gahegan	\$40,000
Grid Middleware Strategy Development	Nick Jones	\$200,000
Community Leadership, Outreach and Governance	Mark Gahegan	\$50,000
Travel for Sustainability, Liaison and Grid Middleware Strategy Development		\$40,000
		\$840,000

The details of these areas of the programme are to be worked through during initial planning sessions to be held with representatives of the community, as outlined in the Milestones section. The bolded lead person for each area will be responsible for leading their members through this process. The related teams will be resourced from the below resource pool, committed by the institutions that have provided support to this proposal.

Table 9: Programme Personnel by Institution and related annual costs

Institution	Person	2009	2010	Total Commitment
Auckland	Mark Gahegan Nick Jones Richard Li Andrey Kharuk Yuriy Halytskyy	267,000	68,400	335,400
GNS Science	Paul Grimwood	20,000	5,300	25,300
Landcare Research	Robert Gibb Niels Hoffman	15,000	5,000	20,000
Lincoln University	Stuart Charters	20,000	5,300	25,300
Massey University	Chris Messom IT Administrator IT Team Manager	87,200	12,800	100,000
New Zealand Supercomputing Centre	Andrew McBeath	20,000	-	20,000
Otago University	Mik Black Tim Molteno	82,000	30,000	112,000
University of Canterbury	Vladimir Mencl	45,000	5,000	50,000
Victoria University Wellington	John Hine Kevin Buckley	82,000	30,000	112,000
Travel budget		32,000	8,000	40,000
		\$670,200	\$169,800	\$840,000



D.3 Expected benefits

ESTABLISHING A GRID MIDDLEWARE STRATEGY

Establishing an national grid middleware strategy is a moajor work area in this proposal, as described above. Our efforts will centre around a core team of ten individuals from seven organisations who will liaise with the wider community of CRI and university stakeholders, as well as other branches of government as appropriate. This team will meet weekly during the project, via videoconference. The outcomes will be:

- A detailed plan for grid middleware buildout (hosted in a Wiki), that is regularly updated to reflect changing technologies, national needs and priorities and opportunities. This plan will be used to direct.
- A dedicated, open-invitation workshop towards the middle of the project where we work intensively on this problem with any stakeholders.
- A strategy document for MoRST, describing the emerging middleware landscape, the opportunities and risks and the strategies we recommend for adopting and sustaining e-research infrastructure (following the workshop).

ESTABLISHING AN EDUCATION PROGRAMME TO ACCELERATE AWARENESS AND ACCEPTANCE OF GRID MIDDLEWARE, GRID SERVICES, GRID TOOLS, WORKFLOW TOOLS, FEDERATED DATABASES.

The Community sustainability section above details plans to ensure appropriate levels of education, with other work areas concentrating on outreach into the science / research community. Education and training outcomes will be as follows:

- Ongoing use of BeSTGRID senior personnel to reach out into their host institutions to raise awareness.
- A sustainability strategy that specifically targets training and education within the development team to ensure that there is a solid foundation of grid-savvy scientists, nationally-distributed.
- Hand-in-hand work with specific groups of scientists in the geo- and bio- realms to construct upper-level eResearch tools and services that these communities can and do use.
- Open access to ALL planning documentation, presentations, training aids and code resources generated during this project.
- An educational workshop specifically targeted at junior researchers and developers, and organised jointly with the Australian ARCS group. This will be co-located with a suitable national conference.

ESTABLISHING OPERATIONAL GRID PARTNERSHIPS WITH OTHER NATIONS

We are already collaborating with Australia (the ARCs initiative, the AuScope geoscience project and Monash University), and will build on these collaborations. Specifically, BeSTGRID is contributing to, and drawing from, Australian efforts to build grid middleware, and we will deepen this relationship, using some of the travel funds requested to ensure New Zealand researchers can attend relevant developer and science workshops.

Our team has ongoing collaboration with researchers at the San Diego Supercomputer Center in the USA, and will add to this some interaction with Argonne National Labs (home of much grid middleware development) and Pacific Northwest National Labs who are leaders in workflow representation tools. Much of the best end-user eResearch products are currently originating from the UK—such as the *Taverna* workflow system and *MyExperiment*. We will build on connections we have with the eScience Centre in Edinburgh, and form a working relationship with the groups at Manchester and Oxford (leaders in the UK). Travel funds will be used where appropriate for these



activities, we will also take opportunities to meet with key individuals and conferences and workshops. Outcomes will be:

- Creation and maintenance of informal social networks to key international institutions and organisations.
- A deeper and broader knowledge of best practice overseas.
- Identification and acquisition of useful middleware infrastructure for translation into the New Zealand context.
- Ability to participate directly in Australian developments to help ensure they also meet our needs.

PROVIDING UPPER-LEVEL SERVICES AND TECHNOLOGIES OF GENERIC USE TO NZ RESEARCHERS, SUCH AS SAKAI OR OTHER ADVANCED WIKIS, COLLECTION MANAGEMENT TOOLS, SOCIAL NETWORKING AND SCIENCE BLOGGING TOOLS, ETC

Please review our workplan (specifically Figure 1 and associated text) which describes our plans from which the following highlights are drawn:

- A broad cross-section of generic grid middleware allowing access to powerful grid computers.
- A level of support for generic research tools and applications including: Wikis, Sakai, Collection Management, job scheduling systems, grid enabled applications such as statistical software (R) and Matlab to help scientists scale up their desktop computing.
- Support for some very specific domain science applications that are of value to targeted research communities (bio- and geo-sciences. These include exchangeable scientific workflows using taverna and Kepler, geovisualisation and mapping web services, phylogenetics tools, access to tools in the Biportal and SCENZ-Grid portal.

An overview of planned services is shown above in Tables 2 and 3 above. More detailed breakdowns of work and deliverables will be developed in early planning sessions, as described in the Milestones section.

ADOPTING APPLICABLE INTERNATIONAL ERESEARCH INFRASTRUCTURES AND TOOLS IN NZ

See above description of grid partnerships with other nations.

PART FUNDING 'NODES' AT ALL CRI'S / UNIVERSITIES TO ALLOW GRID CONNECTIONS.

Funding for node development and maintenance is budgeted for most universities, two CRIS and the NZ Supercomputing Centre (see Table 2). Time and other material support has been budgeted among the development community for technology transfer to other institutions who may wish to become partners with us. Time has also been budgeted for outreach to these organisations to increase awareness and educate. There are current limits to the computational and storage resources available to create nodes at all CRI's and Universities. Those that are able to are contributing these resources inside this proposal. Those that aren't contributing resources won't be left out, as it isn't necessary to create a node at an institution in order that it's researchers access BeSTGRID's services.

D.4 Major milestones

The following table outlines the major project milestones:




Table 10: Activity with reporting and key performance indicators (KPI), organised by year quarters

Area	Milestone	Due	Description	Lead	KPI
Contracting	MoRST Contract signed	Q1 2009	Primary contract between Auckland and MoRST signed	Mark Gahegan	Signed contract
	Subcontracts signed	Q1 2009	Subcontracts with partners signed	Mark Gahegan	Signed contract with each partner
National Grid Middleware Strategy	Bio stakeholder community formed	Q1 2009	Consult with science leaders and establish working relationship with community	Mik Black	Stakeholder committee formed and convened Meeting minutes published
	Geo stakeholder community formed	Q1 2009	Consult with science leaders and establish working relationship with community	Mark Gahegan	Stakeholder committee formed and convened Meeting minutes published
	Bio community priorities established	Q2 2009	Workshop community requirements; prioritise community requirements	Mik Black	Workshop reports published Prioritised Bio community needs published
	Geo community priorities established	Q2 2009	Workshop community requirements; prioritise community requirements	Mark Gahegan	Workshop reports published Prioritised Geo community needs published
	Grid Middleware review of community needs; plan implementation	Q3 2009	Review priorities with science stakeholders Plan implementation of prioritised needs across infrastructure	Nick Jones	Implementation plan published Plan integrated into Grid Deployment plan
Middleware deployment, integration, testing, development	Planning for additional data grid sites	Q2 2009	Workshop requirements with technical working group and additional sites Finalise plan for deployment	Nick Jones	Data Grid Deployment plan published
	Planning for additional compute grid sites	Q2 2009	Workshop requirements with technical working group and additional sites Finalise plan for deployment	Nick Jones	Compute Grid Deployment plan published
	Grid Deployment	Q3, 4 2009 Q1, 2, 2010	Additional sites deployed as per published plan	Nick Jones	Planned sites deployed Publication of newly deployed sites Demonstration of newly deployed sites
Community Sustainability	Infrastructure technical working group established	Q1 2009	Consult with site subcontractor leaders and establish working relationship with technical staff	Nick Jones	Technical Working Group convened Minutes published
	Q2 2009 review	Q2 2009	Review programme of work Plan for knowledge transfer sessions	Nick Jones	Published schedule of knowledge transfer sessions
	Seminar sessions	Q3, 4 2009, Q1, 2 2010	Seminar sessions covering delivered items not yet shared from Q2 & Q3 2009	Nick Jones	Seminar Sessions Presentation materials published



Project Management	Quarterly Reporting	Q1, 2, 3, 4 2009, Q1, 2 2010	Reports on progress to MoRST, Project Management Board, Research Stakeholders, and wider community	Mark Gahegan	Progress report to funder Community update publication
---------------------------	---------------------	------------------------------	--	--------------	---

D.5 Assumptions and risks

We assume that the BeSTGRID partners will remain committed to this project, as they have indicated. Letters of commitment and support have been generated where time allowed². In formulating this proposal, given limited time, we have consulted broadly and reached consensus as a community on the focus of the proposal and the guiding principles. These principles focused on playing to our strengths as a community; demonstrating our successes through strong aligned with research communities; and on building a sustainable base to lay the groundwork for the future. While we have attempted to reach a broad community, limited time for community building and consultation suggests there is a risk we've overlooked some key communities, whom we'd welcome the opportunity to engage with.

Currently BeSTGRID is critically dependent on a few individuals at a couple of sites. To the community to grow to critical mass and achieve sustainability our strategy reduces reliance on key individuals by replicating the skills and services developed. This is achieved by engaging and educating a core community of developers, ambassadors, and leaders, and planning for redundancy in our capabilities (see strategy for Community Sustainability).

Our strategy favours adoption of best-practice software and practices from effective offshore eResearch communities, rather than building our own, to minimise risk and support burdens, and assure compatibility with emerging global infrastructure.

In playing to our strengths, we have focused on research domains where we have direct representation within our Steering Committee. This approach is founded on a need to be literate in the research domains we seek to support, as the translation of researcher needs into effective infrastructure is one of the highest risk and highest value activities we seek to support. We recognise that this approach limits the highest value deliverables to the Geo and Bio related research communities. We would seek to broaden our constituency over time in order that a broader cross section of research communities can be equally represented.

E. Fund management

E.1 Project Managers

The Project Management Team includes the Project Leads and a Project Manager. An experienced project manager, Richard Li will be engaged to manage day-to-day running of the projects, schedules, deadlines, reporting and contracting.

The Project Leads, Mark Gahegan, Nick Jones, John Hine, and Mik Black have many years experience of managing large and multi-institutional research projects, often involving complex software and network related infrastructure. John Hine and Mark Gahegan have many years experience in setting national strategy and vision. All members of the project management team have experience with capability building in education and research sectors.

² Short timeframe for the RFP, paternity leave and sabbatical travel have led to a few omissions that we will take care of in the near future).



E.2 Project Governance

The governance structure for BeSTGRID is already in place in the form of the BeSTGRID Steering Committee. This committee is tasked with governance and strategic leadership, outreach and education. A Technical Working Group is also already functioning within BeSTGRID, it acts as a coordination mechanism across members institutions and will be the key community within which skills transfer and duplication of services occurs.

These groups will be revised to incorporate a stronger project-oriented structure with accountabilities and review processes. This proves difficult in the multi-institutional, collaborative setting that this proposal typifies. Consultation with the supporters and participants within this proposal has led to the following refinements, aimed at achieving overall goals of BeSTGRID as outlined in this proposal.

The Project Management Team will report to the Steering Committee, and seek sign off on plans, reports, education programmes, and communications, and regular consultation on risk and issue registers, and on deliverables. The activities of the Project Management Team will be tracked through existing BeSTGRID systems, including wiki's, document repositories, and collaboration systems. These systems will be maintained by the Project Manager as the authoritative source of project information. The Project Manager will be based at the lead institution, collocated with that institutions Project Leads:

1. The Project Leads will work with their related teams to set detailed plans and to suggest appropriate resourcing strategies. These plans will constitute the significant coordination and control mechanisms within the programme, and will be reported on and reviewed at monthly Steering Committee meetings by the Project Manager, and be the basis for quarterly reports to MoRST and the wider community and stakeholders.
2. The Project Manager has a soft mandate from sub-contracting institutions to work directly with assigned resources and to manage priorities. Due to the difficulties with multiple reporting lines inside and outside institutions, the Project Manager has direct contact with the senior institutional leader to whom escalations can be made.
3. Each senior institutional leader has direct operational management control / authority over the allocated resources, and as a member of the Steering Committee will have full awareness of risks and issues, be responsive to any related to their institution, and take responsibility for completion of deliverables. In the event that any agreed deliverables are not delivered, payments will be deferred until deliverables are complete and signed off by the Steering Committee. This is standard process for sub-contracts between participating institutions, and resulting subcontracts will be constructed to ensure deliverables are explicit and agreed, on a pro-rata basis.

E.3 Project Management Costs

Project management is explicitly budgeted, see budget in Table 3.

F. Motivation

BeSTGRID has dedicated itself to the provision of grid middleware for the New Zealand research community. One measure of its motivation is that it has continued to do this (albeit at a very reduced pace) even after the original TEC funding was exhausted. The rationale for this motivation is that the Universities and CRIs included here see the provision of eResearch as fundamental to their mission. Many organisations have promised additional time, resources and personnel to this project, as a measure of their motivation.

Please see the Letters of Support and Commitment in Appendix and Table 7 above for details of partner motivation and commitment.



G. Conflicts

We are not aware of any conflicts affecting this proposal.

H. Other

Not Applicable

I. Referees

The following Referees have confirmed their willingness to respond to MoRST on this proposal:

PETER W. ARZBERGER

Executive Director, National Partnership for Advanced Computational Infrastructure (NPACI)
9500 Gilman Drive
La Jolla, CA 92093-0534, USA
+1-858-534-5079, Fax: +1-858-534-5056
parzberg@sdsc.edu

IAN FOSTER

Director, Computation Institute
Arthur Holly Compton Distinguished Service Professor of Computer Science
Argonne National Laboratory & University of Chicago
Argonne: MCS/221, 9700 S. Cass Ave, Argonne, IL 60439
Chicago: Rm 405, 5640 S. Ellis Ave, Chicago, IL 60637
+1-630-252-4619
foster@mcs.anl.gov

ANN BORDA

Executive Director, Victorian eResearch Strategic Initiative (VeRSI)
PO Box 4200, University of Melbourne 3052, Australia
ann.borda@versi.edu.au
mobile: 0437 469 417, phone: 03 8344 8322

PAUL BONNINGTON (in case you need a third option)

Professor and Director of the Monash e-Research Centre
Monash e-Research Centre, Monash University
Room W503A, Building 11 (Menzies)
Clayton Campus, Clayton, VIC 3800, Australia
Phone: +61 3 9902 0711, Fax: +61 3 9902 0193, Mobile: +61 4 4753 8018
Email: paul.bonnington@adm.monash.edu.au

J. Contacts

For proposal content and leadership:

Prof. Mark Gahegan

Rm 554, Human Sciences Building,
The University of Auckland,
Private Bag 92019, Auckland

Email: m.gahegan@auckland.ac.nz

Phone: 09 373 7599 Ext 88061

Fax: 09 373 7434

For proposal finances and administration:

Mr. Jonathan Lane

Research Office, 76 Symonds Street
The University of Auckland,
Private Bag 92019, Auckland

Email: j.lane@auckland.ac.nz

Phone: 09 373 7599 Ext 84279

Fax: 09 373 7432



APPENDIX Letters of Support and Commitment

Information and Communication Technology Services

Tel: +64 3 364 2300, Fax: +64 3 364 2332, www.its.canterbury.ac.nz



11 December 2008

To whom it may concern

The University of Canterbury is willing to make available to the BeSTGRID Middleware Proposal the following BeSTGRID resources now housed at the University of Canterbury. Specifically, the GridGateway and underlying virtual machines, the DataGrid with its 20+ terabytes of capacity and access to the BeSTGRID share of the HPC and ECE computer clusters.

Further, the University of Canterbury is willing to make available some capacity of our senior eResearch developer as well as making his time available to work on specific contracted items within the proposal. This time may be used to assist with the integration of the Canterbury resources with new middleware initiatives of this proposal; assisting with the development and installation of new middleware software at Canterbury and elsewhere; training others in these technologies or any other activity that will promote the outcomes from this proposal.

Yours sincerely

Robin Harrington, Advanced Technologies Manager,
Information and Communication Technology Services,
University of Canterbury, Private Bag 4800,
Christchurch, New Zealand.
Phone: +64 3 364 2339 Fax: +64 3 364 2332
Email: robin.harrington@canterbury.ac.nz



11th December 2009

To whom it may concern,

Landcare Research is pleased to fully support the collaborative BeSTGRID Middleware proposal. Landcare Research recognises that the success of and expansion in the use of KAREN is dependant on changes in research modes towards a much more collaborative approach to doing science and that, in turn collaboration will drive success in new science areas.

Landcare Research has already committed strategically to a number of projects that align closely with work packages described in the proposal, and therefore we expect to be able to contribute fully to the outcomes of those work packages. Landcare Research is using the SCENZ-Grid platform to deliver on the FRST funded SpInfo II (Spatial Information) program that supports the LRIS (Land Resource Information System) suite of recognised Nationally Significant databases (1\$m per annum). This applied technology research program is characterised by strong links to a diverse user community in central and regional government. By reinforcing the use SCENZ-Grid within the BeSTGRID community this proposal will significantly empower collaboration and research, and enhance the delivery of end user benefits.

Landcare Research is funding two internal Capability Fund projects that it will contribute (in kind) to the proposal—Semantic-web and Geospatial Visualisation.

- 'Semantic web' is part of Landcare's ongoing support for our internationally recognised capability in Biodiversity Informatics which has contributed to developing the technical infrastructure for federated data-sharing networks (e.g. GBIF, EOL, GISIN), and their underpinning standards and ontologies, e.g. through the TDWG International Biodiversity Standards Group. (\$200k in 2009)
- Geospatial Visualisation is a new capability area and will support SCENZ-Grid (\$364k in 2009)

SpInfo II is promoting an international collaboration to develop SoilsML, which will complement the GeoSciML referred to in the proposal, and extend the collaboration into the international soil modelling community, bringing in links with Australia, USA and EU.

Landcare Research has invested in a geospatial compute cluster (104x cores), a storage site (~26TB), and a virtual collaboration space based on myExperiment. These are core components of our commitment to SCENZ-Grid.

Landcare Research will contribute activities and resources to the project in proportion to our ability to service the needs of the BeSTGRID community. While Landcare Research has made a significant investment in enabling SCENZ-Grid infrastructure, the key factor limiting our ability to contribute at this early stage is our current limited human capacity and the funding to support the expansion of that capacity.

Dr Jerry Cooper
Science Leader – Informatics



Massey University

INFORMATION TECHNOLOGY
SERVICES
Riddet Road
Private Bag 11 222
Palmerston North
New Zealand
T 64 6 350 5163
F 64 6 350 5607
www.massey.ac.nz

To Whom It May Concern

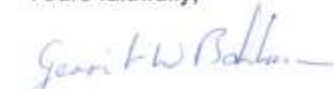
I am writing as the Chief Information Officer of Massey University and as KAREN Champion to give my strongest support toward "BeSTGRID MiddleWare to support New Zealand eResearch", a proposal in response to MoRST's MiddleWare RFP.

Massey University was one of the original BeSTGRID members. Our commitment to BeSTGRID is ongoing with our support of the infrastructure requirements including access to KAREN across the university and providing the infrastructure support for the Massey University hosted BeSTGRID cluster and storage resources. We are committed to extending this support through the next phase of the BeSTGRID Project.

The nature of GRID MiddleWare is to enable cooperation and collaboration between research groups across national and institutional boundaries. Massey University's central IT services has enabled this collaboration with high speed connectivity and secure access to resources via relevant firewall policies. We recognise that the "BeSTGRID MiddleWare to support New Zealand eResearch" Project will further enhance our research users' ability to exploit shared resources both nationally and internationally and we will provide the necessary connectivity and infrastructure support to enable this activity.

Currently the Institute of Information and Mathematical Sciences through the Centre for Parallel Computing and the IIMS IT Support Group manage and maintain the BeSTGRID compute, storage and collaborative facilities. Central IT services have provided the necessary infrastructure support to enable this activity and we are committed to enabling ongoing external access to the BeSTGRID facilities through the next phase of the BeSTGRID Project. We recognise that membership and commitment to BeSTGRID will ensure a world class level of eResearch infrastructure for Massey University and New Zealand's research community and we're looking forward to extending this activity in the project.

Yours faithfully,


Gerrit W Bahlman
Chief Information Officer

9th December 2008



INSTITUTE OF INFORMATION
& MATHEMATICAL SCIENCES
Te Kura Pūtaiao o
Māhiohio me Pāngarau
Private Bag 102 904
North Shore Mail Centre
Auckland 0745
New Zealand
T 64 9 441 8160
F 64 9 411 8136
[Http://iims.massey.ac.nz](http://iims.massey.ac.nz)

08 December 2008

To Whom it May Concern

I am writing as Head of Institute of the Institute of Information and Mathematical Sciences (IIMS) of Massey University to support strongly "BeSTGRID MiddleWare to support New Zealand eResearch" a proposal in response to MoRST's Middleware RFP.


IIMS have many research users of BeSTGRID's High Performance Computing facilities, storage facilities and collaborative infrastructure in Statistics, Mathematics, Computer Science, Information Technology and Astrophysics. Continued access to these and further facilities via the GRID middleware will enable additional and sustained high impact research in these areas.

IIMS has hosted Massey University's activity in the TEC funded BeSTGRID project. As well as the BeSTGRID funded manpower we have invested additional resources to maintain and upgrade the facilities. This has included more than 0.5 EFT of a systems administrator position for the compute and storage resources and additional senior administration and academic staff. As well as the proposed funded manpower contributions in the "BeSTGRID MiddleWare to support New Zealand eResearch", IIMS will contribute an additional 0.6 EFT system administrator and developer as well as an additional 0.4EFT manager/administrator and academic staff to the project. I believe this will further accelerate the implementation of the GRID Middleware in New Zealand leading to significant benefits for the country in the years to come.

IIMS through the Centre for Parallel Computing and the IIMS IT support group have been very active in the implementation and maintenance of the current and past GRID enabled compute, storage and collaborative facilities. Via BestGRID middleware the Massey University hosted BeSTGRID cluster and storage facilities are available for all BeSTGRID members, we are committed to extending this availability through this next phase of BeSTGRID development. We are also committed to maintaining this level of access and support in the future to ensure a world class level of eResearch infrastructure for IIMS, the College of Science, Massey University and of course New Zealand's research community.

Yours faithfully

Prof Tony Norris
Head of Institute of Information and Mathematical Sciences



5 December 2009

Mr Nick Jones
Director, BeSTGRID
Centre for e-Research
University of Auckland
Private Bag 92019
Auckland

Dear Mr. Jones

Victoria University of Wellington sees the development of e-Research as critical to the future participation of New Zealand researchers in projects of international significance.


Victoria has been working on the development of e-Research since 2005 and is pleased to be an active member of the BeSTGRID Steering Group. Victoria has supported e-Research through an e-Research co-ordinator position, partial support of a programmer position, and funding for the development of a computational grid and an OptiPortal. Victoria is also a member of the BlueFern consortium.

Victoria supports BeSTGRID's response to MoRST's RFP concerning Grid Middleware. Should MoRST decide to fund BeSTGRID's proposal Victoria will provide additional support up to 0.5FTE of a programmer position in support of additional work to be done at Victoria.

I wish your proposal success.

Yours sincerely

Professor Neil Quigley
Deputy Vice-Chancellor (Research)



Deputy Vice Chancellor (Research)
PO Box 600, Wellington, New Zealand
Phone +64-4-463 5063 Fax +64-4-463 6798 Email neil.quigley@vuw.ac.nz



11 December 2008

Dr Nick Jones
Manager, Centre for eResearch
The University of Auckland
Private Bag 92019
Auckland

Dear Dr Jones,


The University of Otago strongly supports the proposal put forward by BeSTGRID in response to the recent MoRST RFP for the development of a national grid computing infrastructure for New Zealand.

Otago is committed to the development of eResearch capability, and has a dedicated eResearch Advisory Group (eRAG). This group includes our Director of Information Technology Services, Mike Harte, who is also the Otago KAREN Champion. Otago is a recent member of BeSTGRID, with representation on both the BeSTGRID Steering Committee, and the BeSTGRID Technical Steering Committee.

If the BeSTGRID initiative is successful, the University of Otago will favourably view the proposal to make Dr Mik Black, Dr Tim Molteno and Mr Stewart Stevens available in the FTE proportions specified in the proposal. These individuals will also contribute through their ongoing involvement in the BeSTGRID committees listed above, while Dr Black will contribute an extra 0.1FTE through his involvement with end user groups via the New Zealand Genomics Ltd (NZGL) and Virtual Institute for Statistical Genomics (VISG) projects. In addition, it is anticipated that the bioinformatics component of NZGL will contribute both IT expertise and computer hardware to the BeSTGRID proposal, once the details of this initiative have been finalized.

All the best with your application.

Professor K Geoffrey White
Deputy Vice-Chancellor (Research)



Research & Commercialisation Office

PO Box 94, Lincoln University
Lincoln 7647, Canterbury
New Zealand

Telephone: 64 3 325-2811
Facsimile: 64 3 325-3630

11 December 2008

Professor Mark Gahegan
Director, eResearch Centre
University of Auckland
Private Bag 92019
Auckland

Dear Mark

The University of Auckland ("Auckland") has submitted a bid to the Ministry of Research, Science and Technology ("MoRST") for funding in respect of GRID Computing Middleware Initiative to establish national grid computing infrastructure and middleware tools utilising the Kiwi Advanced Research and Education Network (KAREN) networking capabilities. ("the Programme"). Auckland is to take the role of lead research provider for this Programme. In the event of a successful application, Lincoln University ("Lincoln") will be subcontracted by Auckland to provide services to the Programme.

This letter confirms the agreement made during your recent discussions with Dr Stuart Charters, that the subcontract would involve a minimum of 0.123 FTE (2009) and 0.030 FTE (2010) of Stuart's time in the following work areas, with additional FTE subject to negotiation between Auckland and Lincoln following consultation with the research communities.

Work Areas:-
General Strategy, liaison to other (non-BeSTGRID) communities.

This letter is signed on the understanding that-


1. Lincoln and Auckland each undertake not to disclose any confidential information received from other in the course of developing the Programme detailed in this letter, and agrees not to use any such information for its own benefit without prior consent.
2. Each party shall retain ownership of existing intellectual property they bring to and use for the project. Ownership of intellectual property created in undertaking the Programme and associated matters concerning intellectual property rights and their management shall be the subject of an IP Agreement to be entered into prior to submission of a full proposal.

Signed for and on behalf of Lincoln University:

Date: 11/12/08

.....
Authorised Signatory

Dr Peter John
Director, Innovation & Commercialisation
Lincoln University
.....
Dr Peter John
Director, Research and Commercialisation



Professor Mark Gahegan
School of Geography, Geology and Environmental Science,
University of Auckland, 92019
Auckland 1142.
December 10, 2008.

Letter of Support for Your Proposed 'BeSTGRID MiddleWare to support New Zealand eResearch'

Dear Mark, On behalf of MapNet*, a New Zealand-wide collective of gene mapping scientists, I'm pleased to provide a letter of support for proposal. BeSTGRID is an asset to the NZ scientific community and the developers have enabled our gene mapping community to consider new ways of doing research.

This year MapNet was successful in obtaining a 5-year FRST New Economy Research Fund contract to establish an eResearch-based 'Virtual Institute of Statistical Genetics' (VISG). It involves statisticians and geneticists from the following organisations: ESR Ltd, Scion, Universities of Otago and Auckland, Plant and Food Ltd, ViaLactia BioSciences and AgResearch. This is a new and exciting collaborative model for research in our community of gene mapping scientists. As outlined in our FRST proposal and contract, we plan to use BeSTGRID to host our eResearch portal that is currently under development. This will give our virtual teams a route into NZ-wide grid computing and data storage. Funding of this project will greatly assist our developing collaborations. For example, having a solid, well managed middleware infrastructure that includes training will accelerate our use and uptake both by VISG participants, and the wider MapNet group, with flow-on effects into each of the participating organisations.

The collaborative framework outlined in the proposal is well aligned to MapNet/VISG philosophy of developing underpinning fundamental research that enables collaborations among scientists.

We unequivocally support this proposal, and wish you every success.

Yours sincerely,

Dr Phillip Wilcox
Science Leader, VISG & Convenor, MapNet,
https://mapnet.agresearch.co.nz/mediawiki/index.php/Main_Page

Also:

Senior Scientist
Cellwall Biotechnology Centre
SCION Ltd
Te Papa Tipu Innovation Park, 49 Sala Street
Private Bag 3020, ROTORUA, New Zealand
www.scionresearch.com
phone: +64 7 343 5899 DDI: +64 7 343 5409 cell: + 64 21 387892 fax:+64 7 343 5444
email: phillip.wilcox@scionresearch.com



Gen-i
66 Wydham St
PO Box 8804
Auckland, New Zealand
t: (09) 306 8836
f: (09) 302 1772
e: steve.osborn@gen-i.co.nz
w: www.gen-i.co.nz

10th December 2008

Mr. N. Jones.

Director, BESTGRID
Centre for e-Research
University of Auckland
Private Bag 92019
Auckland

Ref: NZSC Support for BESTGrid proposal

The New Zealand Supercomputer strongly supports BESTGrid's response to MORST's RFP regarding Grid Middleware. The NZSC will provide technical support (in the form of an FTE) in addition to processing and storage capacity to aid the initiative, should BESTGrid be successful with the tender.

The NZSC is committed to developing grid initiatives that support New Zealand's scientific and research communities and was one of the first contracted Service Providers to KAREN over eighteen months ago now. We believe that a research / commercial grid partnership offers a strong platform for New Zealand institutes and welcome both the initiative and the potential for continued involvement with BESTGrid.

Yours sincerely

Steve Osborn
Gen-i Service Line Manager:
Open Source Solutions

A member of the Telecom Group





Te Papa Tipu Innovation Park
49 Sala Street, Private Bag 3020, Rotorua 3046, New Zealand
Telephone +64 7 343 5899 | Facsimile +64 7 348 0952
Email enquiries@scionresearch.com | www.scionresearch.com

SCION 
Next generation biomaterials

11 December 2008

+64 7 343 5441
julia.charity@scionresearch.com

Ministry of Research, Science and Technology
Level 10, 2 The Terrace
PO Box 5336
Wellington 6145

Letter of support for University of Auckland's Research Proposal: "BeSTGRID Middleware to support New Zealand eResearch"

On behalf of Scion, we wish to provide a letter of support for the proposal from the University of Auckland in MoRST's current funding round on GRID computing and middleware to establish a national GRID infrastructure and middleware tools in New Zealand.

The establishment of BeSTGRID in 2006 provided a national GRID infrastructure that has been utilised by geographically disparate research groups around New Zealand. The culture of "eResearch" is gaining momentum and is largely driven off new funding streams supporting cross-sector, multi-disciplinary research teams. It is common-place for many scientific and engineering applications to require access to large amounts of distributed data (terabytes or petabytes). The size and number of these data collections has been growing rapidly in recent years and will continue to grow as new experiments and sensors come on-line, the costs of computation and data storage decrease and performances increase, and new computational science applications are developed.

As outlined in your proposal, BeSTGRID's collaborative framework is one of the underpinning infrastructures that enables collaborations between Scion and researchers globally.

Scion has just recommitted to its KAREN membership until 2015 which provides access to high performance computing (HPC) facilities which enable science projects to be completed faster and cheaper. Without infrastructure like BeSTGRID (and KAREN), we will no longer be able to capitalise on those gains.

Scion has a number of current and future science projects which rely on the availability of BeSTGRID methods and tools to facilitate collaboration on shared information, sharing of computational resources and online visualization of instruments and experiments. For example, MapNet, (a New Zealand-wide collective of gene mapping scientists) are establishing an eResearch-based 'Virtual Institute of Statistical Genetics' (VISG). It involves statisticians and geneticists from the ESR Ltd, Scion (led by), Universities of Otago and Auckland, Plant and Food Ltd, ViaLactia BioSciences and AgResearch. They are currently developing an eResearch portal and intend to host it using BeSTGRID.

We unequivocally support this proposal, and wish you every success.

Yours sincerely



Dr. Julia Charity
KAREN Champion for Scion



To whom it may concern

1 Fairway Drive, Avalon
Lower Hutt 5010
PO Box 30368
Lower Hutt 5040
New Zealand
T +64-4-570 1444
F +64-4-570 4600
www.gns.cri.nz

BeSTGRID Middleware to support New Zealand eResearch

GNS Science recognises that eResearch is becoming an essential part of the general research environment, enabling collaborations with New Zealand and international researchers. We consider the *BeSTGRID Middleware to support New Zealand eResearch* proposal in response to MORST's RFP for GRID Computing Middleware Initiative to be an essential and timely step in the development of eResearch in New Zealand.

We are involved in a number of eResearch projects focussed largely around access to geoscience data:

- 1) a *Seismographic Information Service* providing researchers with fast and machine-readable access to GeoNet's raw earthquake data;
- 2) a *QuakeML* service providing semantically rich access to interpreted earthquake information;
- 3) *SCENZ-GRID*, in collaboration with Landcare Research, providing a platform for wider geoscience collaborative activities;
- 4) the international *One-Geology* project providing a global scale geological map.

We at GNS Science fully support the proposal and wish it every success.

Yours sincerely

Dr Desmond Darby
General Manager Strategy