National eScience Infrastructure (NeSI)

High Performance Computational Platforms and Services for New Zealand's Research Communities

Investment Case

29th October, 2010

THIS IS THE FINAL DRAFT OF THE NeSI INVESTMENT CASE FOR SUBMISSION TO THE MINISTER OF RESEARCH, SCIENCE AND TECHNOLOGY

The Authors

The working group appointed by MoRST and the wider sector to create this proposal comprises:

Prof. Mark Gahegan, U. Auckland (lead author)	m.gahegan@auckland.ac.nz
Mr. Nick Jones, U. Auckland (budget model & finances)	n.jones@auckland.ac.nz
Mr. Rick Christie, independent chair of the working group	rick.christie@xtra.co.nz
Prof. Tim David, U. Canterbury	tim.david@canterbury.ac.nz
Mr. Robert Gibb, Landcare Research	gibbr@landcareresearch.co.nz
Mr. Mike Harte, University of Otago	mike.harte@otago.ac.nz
Dr. Phillip Lindsay, AgResearch	phillip.lindsay@agresearch.co.nz
Dr. Murray Poulter, NIWA	m.poulter@niwa.co.nz

Advisors from MoRST

Dr. Kathleen Logan	Kathleen.Logan@morst.govt.nz
Dr. Wynn Ingram	wynn.ingram@morst.govt.nz

Any suggestions, questions or concerns regarding the proposal can be taken up with any or all of the above persons.

TABLE OF CONTENTS

How to read this document

The document can be read from start to finish, to obtain a rounded picture of all aspects of the infrastructure. An executive summary provides a brief overview of the project. A basic introduction to the current state of affairs and the infrastructure needed to redress the current capability impasse is provided in Section 2. The full science case, including insights into the transformative nature of RS&T priorities is described in Section 3, and the economic case for investment is described in Section 4. Governance details and operating rules and principles are provided in section 5 and budget and financial details in Section 6. Appendices provide some additional details of offshore investments and examples of economic benefits.

It is suggested that all readers examine both Sections 1 and 2 to gain a succinct overview of the current state and the proposal

The Science case (Section 3) can also be read as a stand-alone document, for those interested in the science drivers, impacts and alignment with the goals of research infrastructure provision. This Section is strongly based on the science case prepared for MoRST as part of the Research Infrastructure review process and widely socialised within the RS&T sector.

The Business case, Governance and Budget (Sections 4, 5 and 6) can also be read independently, for those interested in the economic, financial and legal aspects only, and not concerned with the science drivers or impacts.

1 EXECUTIVE SUMMARY

Introduction

High Performance Computing (HPC) and related eScience infrastructure are by now indispensable components of modern science, and are having a major impact on almost every branch of research, throughout the world. This National eScience Infrastructure (NeSI) project will be the most significant infrastructure investment for New Zealand's Science System in the last twenty years. Not only will it enable our scientists across a wide range of communities and disciplines to access and utilise vastly superior computing power, but it also brings the added benefit of achieving this from better and more efficient coordination and cooperation across the sector; rather than proliferating the institutional approach to research infrastructure provision that is common today.

This paper makes the case for government to invest $27.45 million over an initial four-year period towards building a nationally networked, high performance computing and data infrastructure facility. Over the same period, participating members of the research sector will invest $15.41 million in capital assets and contribute $5.56 million towards the operating expenses of this project. Fees for access will also contribute $3.34 million to the budget. At the end of this initial period the project will be evaluated and the basis for ongoing investment agreed between the participating members and government.

This infrastructure will be accessible to all members of the sector who require advanced computing to pursue their research objectives, as well as commercial users. The proposal also provides for the expert support services necessary to allow users to make full and best use of the resources available. The step-change in capability will help keep our best researchers in the country and attract research talent from overseas.

The Economic and Science Benefits

Small countries cannot rely on small science: New Zealand relies on research into some of the world's most complex problems to support its fundamental industries and help fuel a knowledge-led economic recovery. The need is widespread, encompassing every University and CRI. Because of the prohibitive upfront and ongoing costs, no single institution can meet the demands alone; rather it requires collaboration across the research sector to provide the breadth of facilities and capabilities that are urgently needed. There is a major efficiency gain if institutions collaborate, rather than invest separately in the increasingly complex and expensive equipment required. It is more beneficial and cost effective to have this infrastructure created within New Zealand, rather than purchased offshore.

eScience has become for RS&T what eBusiness is for productivity growth in commercial organisations such as *TradeMe* and *Air New Zealand*; and as digital simulations and visualisations have become for the creative industries such as *Weta Digital.*

This investment will accelerate the science that underpins the growth in value-added businesses and directly support the Government's Economic Growth Agenda by enabling more effective, coordinated science and technology-driven outcomes—tied to national science priorities.

As Figure 1 shows, the research enabled by NeSI will cover every 'slice' of the RS&T priorities and specifically will empower many of the nation's major science initiatives such as: the MacDiarmid Institute, High Tech TRST in Materials Technology, Virtual Institute for Statistical Genetics, New Zealand Genomics Ltd, Maurice Wilkins and Allan Wilson health-related Centres, BioEngineering and Bioinformatics, New Zealand Institute for Advanced Study, the Hazards Platform, Greenhouse Gas and Climate Change initiatives, the Liggins and Malaghan Institutes and many others.

The proposed infrastructure will make possible New Zealand’s participation in leading edge research across the globe. In particular, seismic simulation, cross-species gene-mapping and complex biomedical and environmental modelling typically require advanced computing platforms and associated tools and services, which are already available to many other nations. For example, it is critical to New Zealand’s primary industry that we remain at the forefront of technology-based farming and agricultural research and practices. There are huge commercial incentives to scientific success across the life sciences, for example the research-driven ‘Primary Growth Partnerships’ for agriculture, aquaculture, horticulture and forestry. Realistic large-scale modelling is central to many aspects of biology, materials science, geoscience and agriculture across a wide range of contexts.

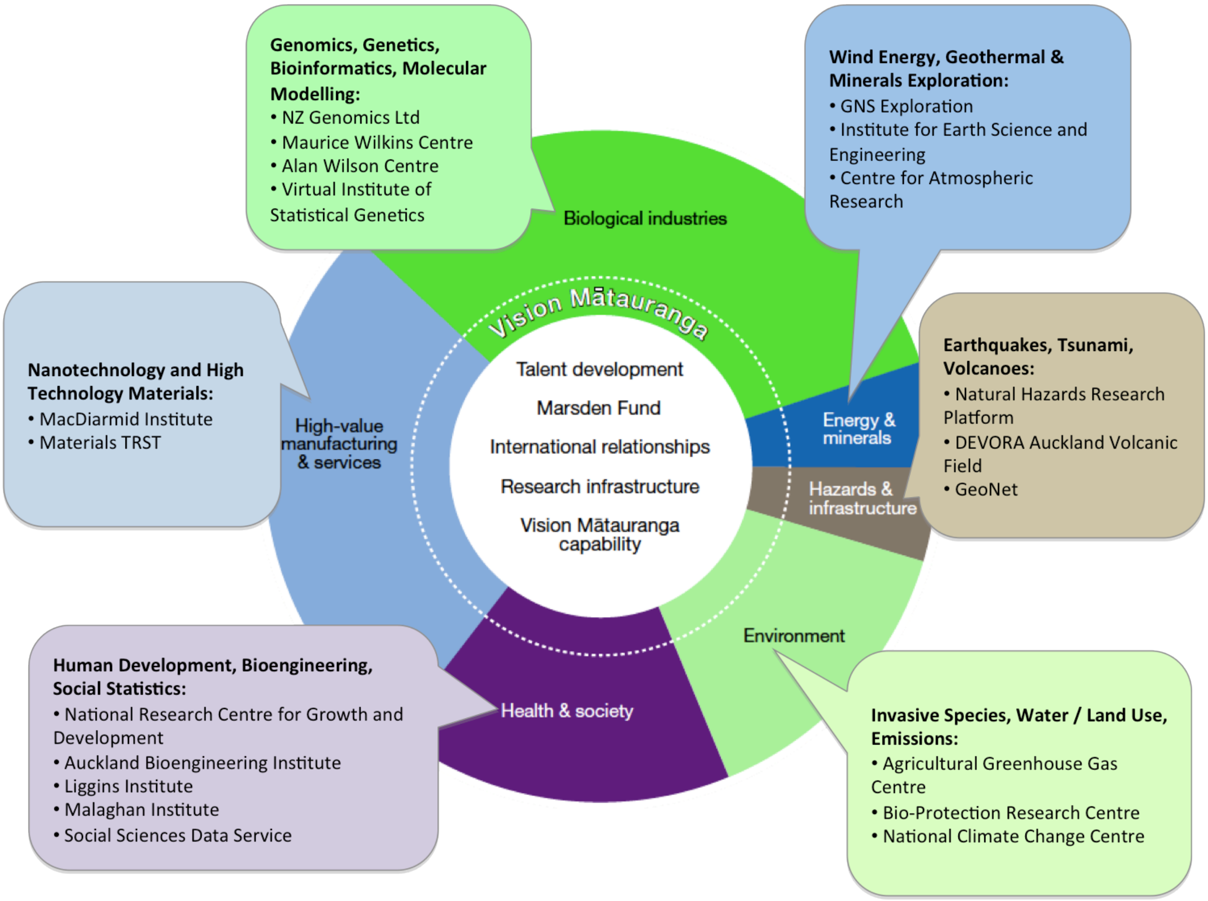

Figure 1. **Alignment with New Zealand’s R, S & T Priorities** showing just a portion of the major science communities that will be served via the proposed infrastructure.

Coordinating National Assets

This infrastructure builds on the existing but limited resources of: the BlueFern HPC facility led by Canterbury, the emerging BeSTGRID eScience community led by Auckland, the recently commissioned NIWA HPC facility and AgResearch’s computational facility. It also includes capabilities at Otago and Landcare Research and connects into several other national-scale initiatives.

The science case associated with this proposal (Section 3) is supported by MoRST, and has gained widespread support among the research sector at large and from many of the country’s top scientists and research groups. Since early 2010, a Transition Working Group (TWG)—selected by stakeholders to represent the RS&T sector and under the guidance of an independent chairman instituted by MoRST—has met regularly to progress the case presented here and to coordinate planning across the sector as broadly as possible. The negotiations have led to a joint understanding of needs, trust as to motives and intentions, and concrete plans to move forward together to create this nationally-focussed infrastructure.

Needs Assessment

A robust analysis of HPC uptake and need showed that existing facilities are either fully utilised, or too costly for researchers to access. Limited access reduces the ability of researchers to conduct modern science. A national survey of computing and eScience needs was conducted, and confirmed these issues. Responses from the survey were used to define the computing platforms, services and capacities required now and anticipated over the next four years initially. These results, which have directly informed our planning and budgeting, indicated the following:

- Current demands for HPC far outstrip capacity, and on some measures by a factor of 25.
- There is a need to provide different types of computational platforms and services, to meet the needs of the nation's advanced science communities.
- Anticipated future demand is many times greater than current demand.
- A shared need exists for standardised and common tools and access mechanisms, which suggests coordinated development of computational services, a 'single front door' approach to coordinating access, and incorporation of security mechanisms, where needed, to protect confidential data.

Governance

The project will be constituted as a legally binding agreement between equal partners, with a single institution as the contracting legal entity. A set of operating principles has been agreed by the partners and is supported by the wider sector. This includes a Charter to capture the institutional agreements covering the operation, provisioning and management of the infrastructure. These agreements will collectively form the foundation on which legal governance will be based. A representative Board, which will include two independent members, will set the policies for transparent and equitable access to the infrastructure, across the science sector as a whole.

Leverage of Existing Research Infrastructure Investments

This infrastructure will provide a major scientific return on significant investments that have been made to date, and will augment the existing infrastructures and human capabilities created so far, and link them via KAREN. There is also an opportunity for New Zealand to access a number of offshore projects currently funded in Australia, the USA, and the European Union, with whom we already collaborate, to leverage their investments here.

Conclusions

New Zealand's research infrastructure stands at a crossroads; the high-speed network is in place, but other vital pieces of a national science infrastructure remain elusive. If these can be provided, New Zealand scientists will be on more of an equal footing with their offshore counterparts, who already have access to sophisticated research infrastructure. Our research communities are united in their need, for and commitment to, a shared computing infrastructure to advance their scientific capability in line with international best practice.

A National eScience Infrastructure (NeSI) will draw together the existing expertise and facilities that are currently fragmented around the country's research institutions and build this into a coherent whole. The result will lift New Zealand's research potential in many fields that are key to economic prosperity and public good.

Government investment is now timely and appropriate to assist the various investing institutions to work together for national good, and to help meet the steep capital and operational costs of HPC equipment and services to be jointly provided in the national interest.

2 INTRODUCTION, AIMS AND PROPOSED INFRASTRUCTURE

2.1 Introduction and Overview

The overall aim of the National eScience Infrastructure (NeSI) is to empower New Zealand researchers to fully realise their potential through the provision of, and ready access to, world-class computational and data-intensive research infrastructure. We already possess abundant talent, drive and imagination that distinguish New Zealand as a prosperous centre for science, engineering, and innovation. Young scientists should regard New Zealand as the locus where they can develop their careers; a place where they may work with the world's best researchers, with the assurance that they will be resourced with a reasonable level of computational infrastructure that will enable them to work and collaborate effectively. Science-led economic growth stands to benefit from a deeper pool of research talent.

Infrastructure requested

Funds are sought to develop a nationally coordinated, state-of-the-art computational facility, and the *eScience services* and research engagement that research communities need. The infrastructure described forms a nationally coordinated ecosystem, building on the pockets of existing expertise and infrastructure available currently. The infrastructure will be initially targeted towards supporting the needs of the nation's advanced science groups, including (but not limited to) those described later in Sections 3 and 4.

Transformative research—by which New Zealand researchers deliver useful outcomes for the country—occurs across a range of different computational platforms and requires various kinds of infrastructure and software support. Specific infrastructure to be acquired or developed is as follows:

1) **Create an advanced, scalable computing infrastructure to support New Zealand's research communities.**
2) **Provide the research tools and applications, data management, user-support and community engagement (i.e. eScience) needed for the best possible uptake and return from this HPC investment.**

NeSI will create an infrastructure that is usable, transparent and scalable, which provides scientists with the correct computing tools for the job. It will further ensure that science teams can work collaboratively, with unimpeded access across the country and throughout the world.

The computing needs of New Zealand's research communities are becoming increasingly sophisticated, requiring access to High Performance Computing (HPC) platforms, vast amounts of data storage, advanced modelling and simulation tools, and increasingly complex experimental environments. HPC is an essential tool of modern science that has a major impact on almost all disciplines. Large-scale, numerically-intensive simulations are providing unprecedented insights and predictive power for global environmental processes, complex biological behaviour, drug design, and the atomic and molecular processes that underlie all of physics, chemistry and materials science. The scope and impact of HPC in research is increasing at a steady pace. Fields traditionally associated with HPC are continuing to flourish, while new communities (particularly in the life sciences) are making large demands of HPC, especially for data mining in the genomics arena, and the modelling of whole biological organs and systems rather than individual cells or system components. Scaling up to address these problems represents a huge computational and scientific challenge, but also provides opportunities for substantial coordination across the research sector.

Table 1 gives useful definitions of the key concepts described here.

Table 1. **Explaining the jargon**: four key eScience terms defined.

Technology	**Description**
NeSI	**The National eScience Infrastructure** that is proposed and described here: a combination of advanced computing, services for researchers and strong research community engagement.
High Performance Computing (HPC) or Supercomputing	**High Performance Computin**g includes state-of-the-art specialised supercomputers and larger clusters of commodity-based processors that allow scientists to scale up their research.
The Grid	**The Grid** is a collection of computational resources often spanning between organisations, coordinated using standardised protocols and interfaces to deliver non-trivial (often sophisticated) services.
Grid Middleware	**Grid Middleware** the set of software services and protocols that integrate the individual computers, storage systems and application tools of research communities into a coherent, distributed, computational environment.
eScience, sometimes also called eResearch or cyber-infrastructure	**eScience** encompasses all of the infrastructure and software that supports a research community. This includes: the HPC platforms, the Grid that connects them and Grid Middleware, along with other services such as, collection management, grid-enabled-science applications, experimental environments, and advanced collaboration tools. Its focus is on deep engagement with researchers to meet their computational needs and to facilitate the formation and operation of effective digital research communities, rather than simply supplying IT services.

What is a National eScience Infrastructure?

National eScience infrastructure is now an essential input to enable the conduct of excellent research. It provides a virtual, distributed, computational and data-intensive research platform geared towards the needs of specific research communities. It allows scientists to focus more on their science and less on the operational details of computation, data management, and shared services. eScience services will allow a researcher to use high performance computing, and online research tools seamlessly and securely from their desktop. Figure 2 shows that eScience Infrastructure is comprised of distributed computing platforms, a high-speed network, layers of software and research communities.

Internationally, such infrastructure is now supporting effective and efficient forms of investigation and collaboration within and between science communities. The development of eScience services has enabled research activities to advance through pooling the resources of individual research groups and institutions to address major research challenges. Appendix 1 describes some of the overseas infrastructures and their levels of investment.

Why do we need it?

The computing needs of New Zealand's research communities are becoming increasingly sophisticated, requiring access to High Performance Computing (HPC) platforms, vast amounts of storage, advanced modelling and simulation tools, and complex experimental environments. HPC is an essential tool of modern science that has a major impact on almost all disciplines.[1] Large-scale, numerically-intensive simulations are providing unprecedented insights and predictive power for global environmental processes, complex biological behaviour, drug design, and the atomic and molecular processes that underlie all of physics, chemistry and materials science.

The scope and impact of HPC in research is increasing at a steady pace. Fields traditionally associated with HPC are continuing to flourish, whilst new communities (particularly in the life sciences) are making large demands of HPC, especially for data mining in the genomics arena, and the modelling of whole biological organs and systems rather than individual cells or system components. Scaling up to address these problems represents a huge computational and scientific challenge; but also provides opportunities for substantial coordination across the research sector.

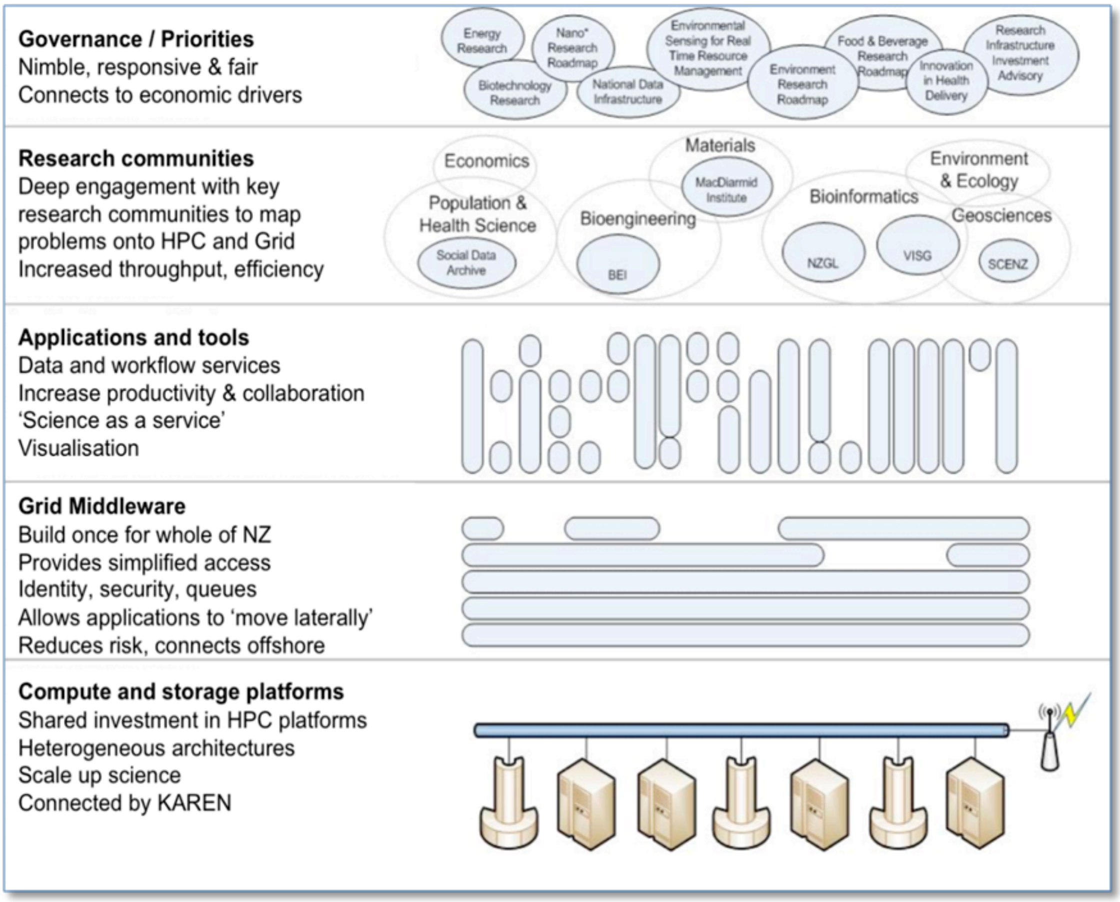

Figure 2. **The eScience layer cake**, connecting the computing platforms to the researchers, via layers of software and services that hide the technical details, making it easier to use.

2.2 What do we Currently Have?

New Zealand's current HPC and eScience infrastructure is in its infancy and has developed on an *ad hoc* basis nationally. It has wide support and has grown from sector demand but is under-funded and over-utilised. There are pockets of excellence, but not enough capacity to meet the country's needs. Most existing infrastructure has no longer-term funding strategy in place, so is only able to offer short-term service guarantees to its user communities. The key challenges and barriers to growth of eScience in New Zealand include:

---

[1] See, for example: *Scan of New Zealand's Large Scale Research Infrastructure Needs*, 2007-2012, Report by the Research Infrastructure Advisory Group, available at: http://www.morst.govt.nz/current-work/science-infrastructure/research-infrastructure/scan-2007-2012/

- A dire shortage of capacity, and of capability.
- Lack of critical mass and issues of inefficiency due to a dispersed skill base and management of advanced ICT infrastructure.
- No longer-term funding in place to foster trust and encourage uptake.
- Inadequate resources for participating in global collaborations.

These barriers have discouraged scientists from exploring research problems that require significant computational resources; as a consequence, many scientists restrict problem-solving to that which is manageable on desktop systems, thus contributing to a growing gap in New Zealand's science and technology capability relative to offshore. Internationally, such capability is being developed to support multi-disciplinary research programmes that meet strategic science and technology priorities, working across institutional and international boundaries.

What is needed to improve it?

Financial investment is urgently needed to develop a nationally-coordinated approach to provide state-of-the-art computational hardware platforms (i.e. HPC and computational clusters), grid middleware and specific application software, along with engagement with science communities to leverage these resources to maximum effect.

Developing this infrastructure will not only facilitate better science outcomes in key fields, but also lead to significant improvements in procurement efficiency. Individually, New Zealand institutions cannot sustain eScience incremental growth within their current research budgets, let alone support strategic leadership in these areas. Rather than fund such initiatives piecemeal by institution, as is currently the case, funds would be far more effectively used to create and implement a national Science infrastructure base to support New Zealand's research needs. A similar strategy has been adopted in a number of countries where the development of major eScience infrastructures is now in an advanced state. This is the case throughout Europe, Australia, USA, as well as China, India, and South Korea.

2.3 Market and Coordination Failures and Opportunities

The market conditions that prevent sharing of HPC and eScience infrastructure result in lack of coordinated investment. Furthermore, the added costs of coordination result in failure of co-investment, even where efficiencies may be gained. There is significant opportunity for system-wide efficiencies to be gained by coordinated infrastructure investment and service provision.

The current system, where there is no nationally co-ordinated eScience infrastructure, has a number of shortcomings, which include:

- Competition between individual institutions discourages co-operative or shared infrastructure.
- There is some duplication of equipment between research groups and institutions. This can be costly and inefficient; not only do institutions (and the sector) lose the benefits of bulk purchasing power, but also researcher hours are wasted in managing fragmented computing systems, at the expense of research productivity.
- The lack of a secure committed funding stream reduces the impact of current programmes, and discourages their uptake. Because ongoing investment is either short-term or undefined, individual institutions are forced to duplicate equipment to ensure that their own research programmes are not jeopardised.
- At present there are no guarantees that much of the current infrastructure (e.g. BeSTGRID's eScience program and Canterbury's BlueFern HPC facility) will exist beyond 2010, because the market cannot sustain them without additional investment.

- Traditionally, research grants do not fund dedicated resources for computationally–intensive research. Yet to remain internationally competitive, many researchers need to be using HPC.
- The high costs of ownership and high rate of depreciation that apply to advanced computational platforms prevent institutions from providing free national access to HPC facilities. Yet, cost-recovery user charges create a price barrier to most researchers.

Enhancing collaboration

The CRI task force report (*How to enhance the value of New Zealand's investment in Crown Research Institutes*) recommends that the government support inter-institutional collaborative research. This infrastructure will provide a key platform in support of increased collaboration, which fosters joint development and provision of research computing tools, services and platforms across the sector. It would specifically encourage:

- Joint development and operation of expensive research support services such as: operating computing platforms, job submission and identity management.
- The creation of new research data, and the infrastructure for accessing and sharing the data more effectively.
- Support for virtual research communities, collaborative environments and access to remote instruments.

In addition to encouraging and facilitating new collaborations, it will also enhance the research activities of several collaborative consortia that require sophisticated computational platforms and services. These consortia have been established in direct response to New Zealand needs and are guided by strategies with significant stakeholder input and strategic advisory groups. For example, NZ Genomics Limited, the Natural Hazards Research Platform, the New Zealand Agricultural Greenhouse Gas Research Centre, the MacDiarmid Institute, the Maurice Wilkins Centre and many others are established research collaborations that would benefit from the mature, stable services and expertise that NeSI will deliver. This includes opportunities for better collaboration on infrastructure procurement and provision, and the opportunity to purchase reliable computational capacity for their unmet research needs. We will target the above, existing collaborations in the first instance.

Finally, collaboration with overseas research groups and infrastructures will also be greatly enhanced, enabling NZ-based researchers to be directly integrated into many of the existing eScience communities that have formed (or are forming) internationally. The infrastructure we create will have many points of connection above the network to a growing number of nations who have invested in national-level eScience.

2.4 Forming a National eScience Infrastructure

The National eScience Infrastructure will be formed via legally binding agreement between the four partners (Auckland, Canterbury, NIWA and AgResearch), supplemented by Otago and Landcare Research. Each of these partner institutions has pledged to commit significant co-investment to match the requested crown funding, and to abide by a charter and related governance and operating principles—all fully described in Section 5 and Appendix 5. The partnership brings together the country's major centres of excellence in HPC and eScience. Figure 3 below depicts this partnership.

Capacity will be created within this partnership to service not only local institutional needs but also those of additional strategic research communities. NeSI will employ a small team of specialist support staff to operate the infrastructure, help users access the facilities effectively and create the eScience services that provide the ease-of-use and desktop functionality that researchers need.

Access to the HPC platforms will be provided by several mechanisms including:

1. Host institutions will gain access to capacity roughly equivalent to their investment.
2. The research sector at large will be offered access by merit to the remaining capacity. Much of this capacity will be used to support publicly funded research, contributing a proportion of the cost (subsidised access). Planned charges will be 20% of the actual costs.
3. An emerging research and researcher fund will be offered to provide a small amount of free access to emerging researchers or priority areas, by setting aside some capacity.
4. Commercial users will pay the full costs to use the infrastructure.

Note that all merit users will, where practical, access the infrastructure through a 'single front door'—that is a common set of interfaces to hide operation details from the researchers. Additional details of governance, and rights to access are provided later in Section 5.

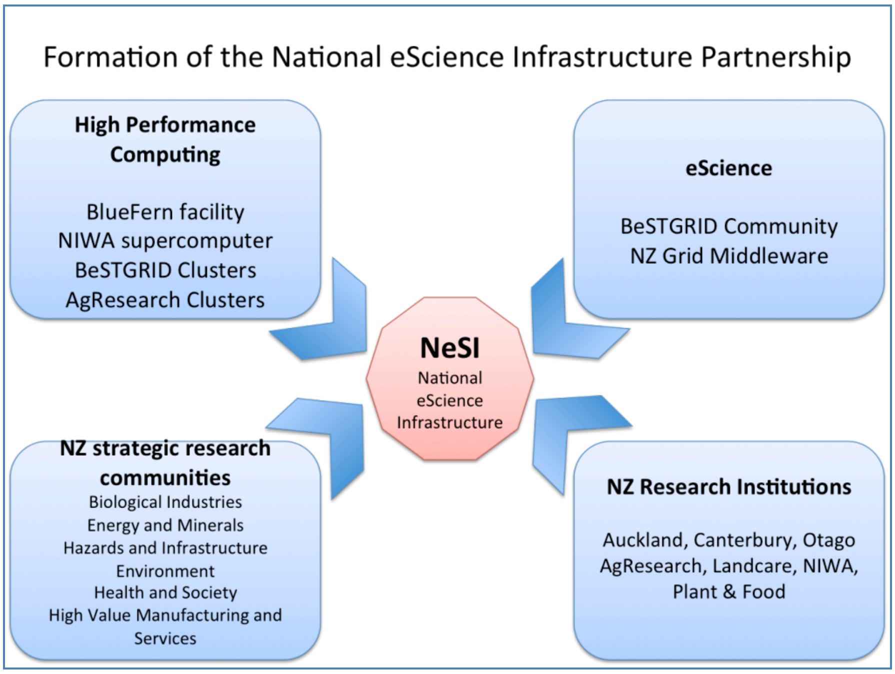

Figure 3. **The NeSI partnership**: a collaboration between high performance computing facilities, eScience infrastructure, research communities and research institutions.

2.5 The Principal Institutions and their Capabilities

The current infrastructure and capabilities at the principal institutions, described below, comprising the Blue Fern HPC group, BeSTGRID, the NIWA High Performance Computing Facility (HPCF) and the HPC facilities at AgResearch[2]. While this infrastructure has developed to date based largely on individual institutional requirements, there has been some technical collaboration recently between institutions through the BeSTGRID alliance.

With the exception of NIWA's new HPCF, the current infrastructure is fully utilised.

A summary of the infrastructure and capabilities of each of the principal partners follows. A fuller description of each partner can be found in Appendix 2.

2.5.1 BlueFern HPC Facility

BlueFern®, a collegial, HPC and eResearch services facility based at the University of Canterbury, is a key component of the current NZ HPC and eScience environment. It has been in existence since mid-

---

[2] www.bestgrid.org/index.php/Main_Page; www.bluefern.canterbury.ac.nz; http://www.niwa.co.nz/our-services/hpcf; http://www.agresearch.co.nz/science/bioinformatics/capabilities.aspx.

2006, beginning with an IBM p575 based computer system and augmented in mid-2007 by the addition of the first IBM BlueGene in the southern hemisphere. Currently, the BlueFern infrastructure consists of two IBM supercomputers, a P-Series and a BlueGene, housed in a secure, reliable data centre at the University of Canterbury. It supports a large and varied community, with a strong external user-base of researchers from around the country.

2.5.2 BeSTGRID eScience and Grid Computing Capabilities

Led by the University of Auckland, BeSTGRID is New Zealand's national eScience community, providing services and infrastructure to support research computing. BeSTGRID coordinates shared eScience infrastructure across a growing partner network that includes the universities of: Auckland, AUT, Canterbury, Lincoln, Massey, Otago, Victoria, and Waikato; and from the CRI sector: AgResearch, IRL (MacDiarmid) and Landcare. BeSTGRID has in place experienced, skilled, software development and support, and managerial staff at key NZ research institutions, along with strong engagement from domain scientists. It is funded through to the end of 2010 via a MoRST initiative.

2.5.3 NIWA High Performance Computing Facility

NIWA has had HPC capability since 1999, for use in water and atmospheric modelling to support public good research in New Zealand and international programmes such as IPCC. Due to NIWA's time-critical needs, it maintains an ownership model based on facility management at NIWA. NIWA has recently invested in an upgrade of their HPC infrastructure; the new supercomputing and data management facilities are located at NIWA's Wellington site at Greta Point. With appropriate storage and communications, the facility has been scoped to meet demand for 5 years with a mid-life upgrade planned to double the capacity. Access will be provided through the infrastructure proposed here.

2.5.4 AgResearch Computational GRID

AgResearch has been a leader in high performance computing in New Zealand with its initial foray into the field in 2000. Through this period the AgResearch facility and its staff have maintained AgResearch's position at the leading edge of molecular and bioinformatics based research, including large-scale microarray design, manufacture, experimentation and analysis, and proteomics analysis. More recently, from 2006 to the present, this facility was the main resource underpinning assembly of the first drafts of the sheep, deer and fungal endophyte genomes, and will over the next five years lead the data processing and bioinformatics required for assembly of further pastoral forage plant species, and increasingly, individual genomes of some thousands of genetically valuable livestock.

2.6 Research Community Consultation and Needs Assessment

This proposal is founded on four years' experience operating BeSTGRID and BlueFern services, and ten years of NIWA and AgResearch HPC purchasing and operating experience. It is informed by needs expressed by stakeholders during ongoing community engagement within these arenas and via feedback from the Research Infrastructure Advisory Group (RIAG). The proposal has been further developed through:

1. consultation with nationally significant research communities who support the science case, as described in Section 3,
2. a national survey of HPC and eScience needs, conducted in April 2010, to ascertain both the type and the scale of the required infrastructure (described in the following sub-section), and
3. detailed discussions about institutional research IT needs with representatives (CIO and IT Director level) from all participating institutions.

To identify institutional needs, several organisations have convened internal working groups to assess their scientists' demand for eScience infrastructure, and are evolving institutional strategies and organisational units (often established as Centres for eResearch and HPC) to execute them, with advice from a group of science stakeholders. The institutional role of coordinating science user IT needs, including quantifying differing requirements for HPC, large-scale data storage, and applications and services infrastructure, rests with these working groups. In a few cases, the working groups have developed service delivery models, including forecasts and costings of growth in demand over the coming 5 years. Complimenting this is their institutional knowledge about the scale of matching human resources needed to effectively implement, maintain, and support the eScience infrastructure services.

The scale of this investment was determined via MoRST's negotiations with government, and the matching capital that the major partners could bring to bear. Review of current needs and extrapolation into the future suggest that the investment will not meet all demand, but it will provide for a large portion of it. Much work has taken place during the planning phase to assess the type and proportion of computing architecture needs across the different science sectors. The intention is to ensure a mix of computing resources of appropriate architectures that efficiently meets the communities' needs.

2.7 Investigation and Survey of HPC and eScience Needs

April 2010 saw a national survey of HPC and eScience needs conducted. Respondents were asked to assess several aspects of importance here, including: the preferred type of HPC platform—both now and in the future, current and desired level of access, tools and services judged to be vital, commonly used applications and packages, and many other aspects. Approximately 200 responses were received, with many of these being pooled, where a community or laboratory provided a single summary response.

The major findings are as follows:

- Currently, there is a serious shortage of HPC capacity throughout New Zealand, and this hampers research productivity. This surfeit will become much more severe over the next five years.
- There is widespread need and support for several HPC platforms, including super-clusters, traditional HPC platforms and specifically high memory and high processor count architectures.
- Newer HPC technologies such as General Purpose Graphical Processing Units (GP-GPU) are needed now and will become highly sought-after in the following five years.
- Researchers clearly see the need for eScience services, such as: Identity Management, access to storage infrastructure and Grid Middleware.
- Researchers greatly value simple interfaces that hide the complexities of accessing remote computing facilities.
- There are sets of applications and productivity tools that are common between many researchers, suggesting that central provision and support would be improve efficiency.

These findings were used to inform the balance of investment across computing platforms and eScience infrastructure, as described in the Budget Section. Further results from the survey are shown in Appendix 4.

2.8 Infrastructure Required

Below are details of the components needed to develop a nationally coordinated, state-of-the-art computational infrastructure, including the grid middleware and specific application software that research communities need.

Transformative research, by which New Zealand researchers deliver useful outcomes for the country, occurs across a range of different computational platforms and requires various kinds of infrastructure and software support, informed by the needs of users. Hence the plans below provide a mix of computing platforms and support infrastructure, in direct response to the clear needs described in the previous sub-section.

Specific infrastructure to be provided is as follows:

1. **Create an advanced, scalable computing infrastructure to support New Zealand's research communities:**
    a. **Provide state-of-the-art High Performance Computing platforms,** to support a wide range of applications. The facilities will be specified and scaled to fit national needs, and will comprise: a high-processor count, low power HPC platform and a high speed, high memory HPC platform. These platforms could be linked to provide a multidisciplinary, shared, peak computing capability for the most computationally-demanding scientific research as well as for experimental super-computing.
    b. **Create two national super-cluster facilities** offering a variety of differently configured compute nodes[3] (i.e. small collections of CPU cores) that will give flexibility in tackling a broad range of computational problems and additional capacity for handling large, but highly parallel, simulations. Some nodes will be in the form of Graphical Processing Units (GPUs).
    c. **Encourage the growth of grid-connected compute clusters already in place and integrate additional clusters at other universities and CRIs** to create a nationally distributed, coordinated, heterogeneous grid infrastructure. This would also cover commodity cloud computing along with campus clusters such as those provided by lab-based computing pools and other emergent HPC technologies as and when they become useful.
    d. **Manage and operate these facilities in a coordinated manner** to address systems administration, scheduling, software licenses, access control, and migration of user-applications and data around the infrastructure as needed.
    e. **Create a distributed, compute platform** to host the middleware, applications and services required by research communities, thus providing a common, shared infrastructure across the country, with a uniform means for researchers to access it from anywhere.

These computational resources will match the diversity of needs for HPC services, enabling seamless access to two research-focused supercomputing resources at NIWA and Canterbury, super-clusters at Auckland and AgResearch, several local clusters and associated storage. Details of the expected specification of these platforms are provided later in Section 6. Specifically Table 5 shows the budget breakdown for capital expenditure and investment and Table 6 lists indicative capabilities that we expect to obtain for this expenditure.

2. **Provide the grid middleware, research tools and applications, data management, user-support, and community engagement needed for the best possible uptake and return from this HPC investment:**
    a. **Support grid middleware across the above HPC infrastructure.** Ensure seamless access to HPC resources for New Zealand researchers by enhancing the established and successful

---

[3] For example, nodes with faster processors, or larger local memory, or faster interconnect.

BeSTGRID middleware community. A national development team will focus on middleware build-out, generic research services, federated access, and associated data storage and management systems, using overseas best practice.

b. **Support specific engagement with nationally-strategic research communities.** Working with science communities, migrate the specific research tools and applications needed into the HPC and grid middleware environment. Add other productivity tools such as repositories, data integration services, workflows and service chaining, and sophisticated *in-silico*[4] modelling software. **Help migrate established research communities and individuals, as necessary,** to these new technologies to ensure their successful adoption[5].

c. **Manage the continued development of this nationally-coordinated eResearch infrastructure,** via strategy development and technology transfer from relevant organisations offshore. Where we are able, **play a significant role in international cyber-infrastructures and eResearch projects** to ensure that our infrastructure "connects to the world".

d. **Create a consulting service to help researchers understand what NeSI resources they need, and how to access them.** This service would ensure smooth access to all platforms and services, with appropriate levels of help to users.

3. Provide a nationally-distributed storage infrastructure to support the advanced needs of computationally-intensive research:

a. **Develop and provision grid-connected, mass storage infrastructure** required to support not only key science databases and collections but also the data-intensive applications and simulations that will run on the above platforms. Three classes of storage (3-tier) will be supported, depending on data volume, access patterns, security, reliability and access speed. Storage will be distributed around the infrastructure to reduce risk and improve performance.

b. **Storage will be distributed around the infrastructure to reduce risk and improve performance**. It will be scaled to meet the science applications described below, with storage nodes placed close to the computational centres or important scientific equipment, fully connected by the infrastructure. The tiered storage will be available as a national resource and will support controlled, secure access from remote locations, via the grid middleware.

c. **Significant annual growth in capacity will be planned to meet rising demands**. Storage infrastructure will lay important foundations for sharing nationally significant datasets and for future building of a national, shared data fabric.

Massive storage capacity is critical to underpin all HPC activities, be they data driven science or in-silico simulations. For the national system, storage will be provisioned in tiers. A typical 3-tier model provides the flexibility needed to store large datasets for immediate, fast access e.g. the Bioinformatics (BioMirror) and ocean biodiversity (Obis) datasets hosted by BeSTGRID, through to archives of experiments that require slower, occasional access.

Section 6 contains further details of the budget, staffing levels and capabilities for the above facilities and services.

This proposal will create an infrastructure that is usable, transparent, scalable and provides scientists with the correct computing tools for the job. It will further ensure that research teams can work

---

[4] Many research communities now use very sophisticated simulation environments for modelling and analysis (subjects as diverse as climate, plate tectonics, biomechanics, material science and land cover change). *In silico* refers to experiments performed on a computer.

[5] The support to transition people onto HPC platforms, is in addition to support for existing users. A third area of support will require development of training opportunities and materials, which NeSI will also address as part of its core business.

collaboratively, with unimpeded access across the country and deeper connections to offshore science infrastructure.

2.9 Shared Investment into a Small Number of Centres of Capability

Rather than distribute the above investment right across the sector, centres of capability (Section 2.5) will be developed at a small number of locations, spread across the country (Auckland, Wellington, Christchurch, Dunedin). The geographical location of facilities will not affect their access or uptake, but will allow for efficiency gains in purchasing, housing and operating.

3 THE SCIENCE CASE

The science concept proposal for this investment was approved by the RIAG advisory committee in late 2009, and presented to representatives from all of the public research institutions in February 2010, where it received unanimous support and endorsement. The main elements of the science case are presented in this section.

3.1 New Zealand Research Priorities

New Zealand's research, science and technology (RS&T) priority areas are high-value manufacturing, biological industries (primary industries and food), energy and minerals, hazards and infrastructure, environment, and health and society, as shown in Figure 4 (research priorities as depicted in the recent government document: 'Igniting Potential: New Zealand's Science and Innovation Pathway'). There is a very strong alignment between our proposed support of research communities and these government RS&T priorities, described below in Section 3.3.

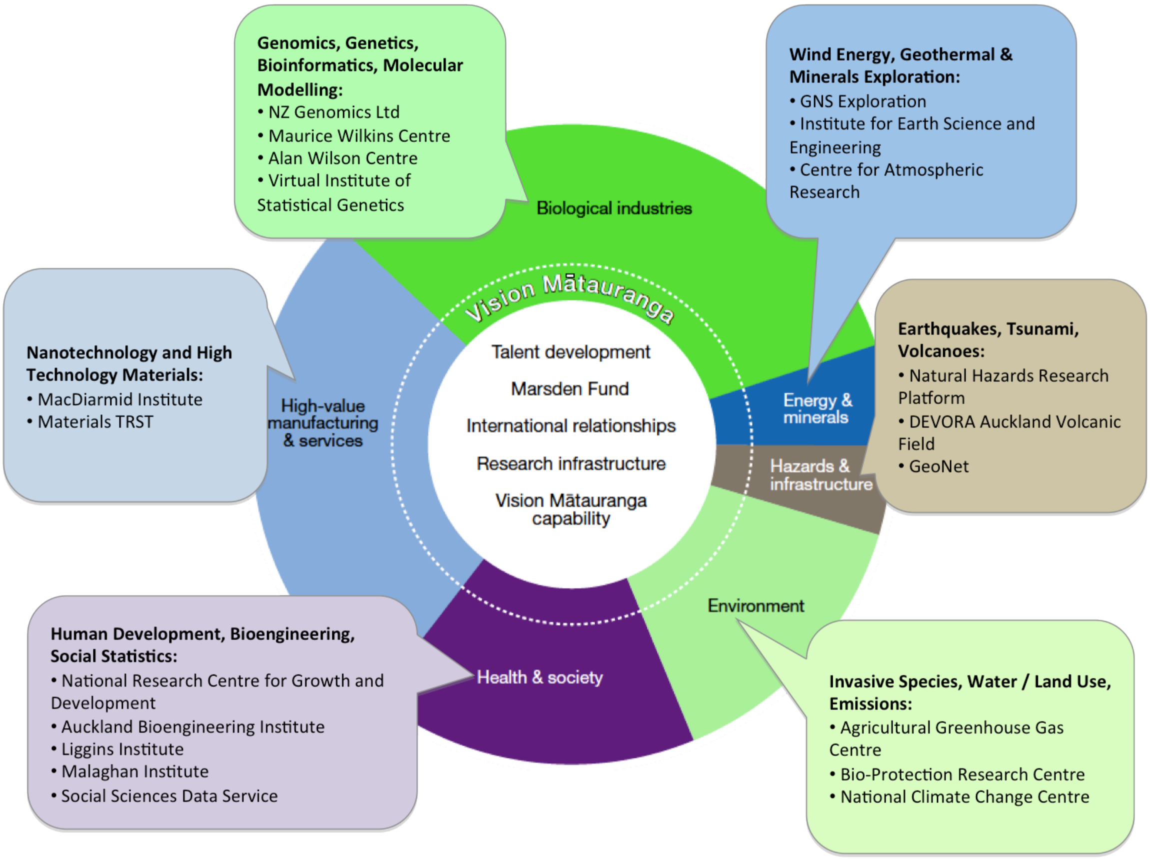

Figure 4: **New Zealand R, S & T Priorities**, with associated research communities that are or could be, major users of an eScience Infrastructure. This figure also appears in the executive summary.

3.2 Vision

Internationally, eScience—the infrastructure to provide a virtual, distributed computational and data intensive research platform geared towards the needs of specific research communities—is now supporting effective and efficient forms of investigation and collaboration within and between science communities. It has enabled research activities that integrate the limited resources of individual research groups and institutions to address some of the bigger research challenges, allowing scientists to focus more on their science, and less on the operational details of computation, data management, shared services and collaboration. Rather than fund such initiatives independently as is currently the case, funds would be better used to **create and implement a national eScience infrastructure, to support New Zealand's research needs in this realm**—a strategy that is being used to good effect in many other countries. The development of major eScience infrastructures is now in an advanced state across all of Europe, Australia, the USA, Canada, Japan and even into China, India and South Korea.

HPC includes state-of-the-art specialised supercomputers and larger clusters of commodity-based processors and data storage. eScience encompasses all of the infrastructure and software that supports a research community using these sophisticated computing platforms and includes: the grid middleware[6], federated identity management, storage management, grid-enabled-science applications and experimental environments and advanced collaboration tools.

3.3 Alignment with Government Policies for Research

New Zealand has taken a major step towards 21st century scientific infrastructure in developing the Kiwi Advanced Research and Education Network (KAREN), and some of the more basic associated eResearch technologies such as advanced video conferencing, identity and access management and grid computing[7]. By connecting computing and storage resources at different locations, via KAREN, into a seamless computational grid, and building on these existing programmes, this proposal will provide a major scientific return on the investment that has been made so far, as well as providing stronger arguments to sustain the infrastructure and human capability thus far created.

MoRST's *2010 Research Infrastructure Strategic Plan* identifies HPC and eScience services as the highest priority, broad-impact research infrastructure need for New Zealand. This sits alongside, and requires the high-speed broadband research and education network to function optimally.

The proposal also aligns well with Government policy for research infrastructure, as expressed in the four MoRST roadmaps for science and it aligns with the Transformational Research, Science and Technology (TRST) areas. For example, Section 2.3 on Global Research Trends reports on the:

> *"…growth of numerical models as tools for investigating processes; availability of more sophisticated mathematical and statistical tools; adoption of new methods for studying complex adaptive systems; and a growing ability to blend different models and datasets."*

In Section 4.2 under Systems Understanding and Integration:

> *"...modelling that links different programmes and disciplines, ranging from conceptual models to detailed computer models; scaling up science; multidisciplinary capabilities.* "

The proposed national framework is also strongly aligned with the concept of research "platforms", as presented in the recently-circulated document "*New Zealand's Research Science and Technology Priorities*".[8]

By catering for the expanding computing and eResearch needs of the NZ research sector, the proposed framework will provide the underpinning infrastructure on which new strategic research platforms can by built, allowing strategically funded research in areas of national importance to proceed without duplication of effort across distinct platforms.

---

[6] See definitions of these terms provided earlier in Table 1, Section 2.1.

[7] http://www.morst.govt.nz/current-work/science-infrastructure/eresearch/

[8] http://www.morst.govt.nz/current-work/NZ-RST-priorities-feedback/

This reflects the third point of the "Operational Principles" outlined in the RS&T Priorities document (page 6):

> *"To foster efficiency, emphasis should be given to where a multi-organisational approach is possible so that critical mass can be achieved, duplication is avoided, advanced infrastructure can be developed, and latent and real synergies across partners can be exploited*."

Finally, the collaborative aspects of eScience will be a vital component in building the international linkages and relationships that are highlighted throughout the document.

Support for ongoing and planned national research initiatives

Some large-scale science and technology projects planned or underway will depend on this infrastructure to provide the necessary tools, services and compute power. These include the Statistical Genetics and NZ Genomics groups.

While this infrastructure does not have sufficient funds to directly support the Square Kilometre Array (SKA) radio-telescope, any future involvement of NZ in this landmark international research infrastructure will require significant hardware and associated human capabilities to be well established in New Zealand. Science-driven, post-processing opportunities exist for NZ researchers, if they are suitably equipped and trained.

Additional computationally-based infrastructures, such as a national data fabric (coordination, sharing and broad access to the country's research and cultural data assets) will also benefit from the coordination of computational infrastructure and services.

A nationally connected HPC and eScience infrastructure will dovetail with the requirements of these and many other research investments, optimising outcomes.

3.4 Details of Support for the Nation's RS&T Priorities

HPC is an essential tool of modern science that has a major impact on almost all disciplines[9]. Large scale, numerically intensive simulations are providing unprecedented insights and predictive power for global environmental processes, complex biological behaviour, drug design, and the atomic and molecular processes that underlie all of physics, chemistry and materials science. There exists no slackening of the pace at which the scope and impact of HPC in research is increasing. Fields traditionally associated with HPC are continuing to flourish, while new communities (particularly in the life sciences) are making big demands of HPC, especially for nano-technological devices, data mining in the genomics arena and the modelling of whole biological organs and systems rather than individual cells or system components.

Scaling up to these problems represents a huge computational and scientific challenge. Progress is being made in nanotechnology with fundamental work by the MacDiarmid Institute, and in bioinformatics and bioengineering at several universities and CRIs. Indeed, a common theme that runs through many of the scientific areas is the integration of several computational techniques to attack problems with a number of different spatial and temporal scales. As well as being ***multi-scale*** they are often also ***multi-disciplinary***—involving several teams and specialisations. Systematic methods for breaking these scales into manageable but meaningful chunks are under investigation everywhere[10].

---

[9] For example: See *Scan of New Zealand's Large Scale Research Infrastructure Needs*, 2007 – 2012, Report by the Research Infrastructure Advisory Group – available at http://www.morst.govt.nz/current-work/science-infrastructure/research-infrastructure/scan-2007-2012/

[10] For example, see the *BioComplexity* program in the USA: http://www.nsf.gov/geo/ere/ereweb/fund-biocomplex.cfm.

The following descriptions of the nation's key RS&T challenges follows the same structure as shown in Figure 4 above.

3.4.1 Biological Industries

The arenas of genomics, biomolecular and bioinformatics is of key scientific interest to many NZ institutions, including: AgResearch, Otago, Scion (Virtual Institute of Statistical Genomics), the Auckland Bioinformatics Institute, Maurice Wilkins Centre for Molecular Discovery and NZ Genomics[11]. These organizations require considerable HPC support in terms of data mining and gene mapping. It is crucial to New Zealand's fundamental industry that we continue to remain at the forefront of DNA technology-based farming and agriculture.

In biomolecular simulations, realistic large-scale modelling is central to many aspects of current biology in a wide range of contexts. Ongoing achievements in structural and functional genomics challenge simulations to provide an integrative view of the dynamics of multi-component biological systems, whether at the level of protein molecules, of macromolecular assemblies or of cells and organs. The demands placed on HPC by these tasks are massive, requiring dedicated use of large-scale, high processor count computing over extended periods of time. Thus, biological simulations will be extended beyond their more traditional domain of small protein molecules to analyse more complex, multi-subunit proteins, at complexes of multiple proteins and larger scale assemblies within cells, and importantly at modelling of complex tissues and organs such as the heart, kidney, liver and brain. There is considerable work to be done at the interface between bioelectrical simulation, biomedical imaging and clinical neuroscience that will place considerable demands on HPC in terms of simulation capabilities, support for very complex suites of software and truly massive storage needs. This leads us towards two of the grand challenges for biological simulation, namely computational neuroscience / brain simulation studies and whole-organ modelling.

The National Research Centre for Growth and Development in its remit states that the overall goal of their work is "to better understand the effects of environmental factors acting during mammalian development on a range of intermediate and long term outcomes relevant to human health and agricultural production". Although usage of HPC does not immediately spring to mind, the calculation of environmental factors affecting mammalian growth (especially with the young) is becoming a significant computing challenge with epidemiological factors a major contributor in the later onset of cardiovascular diseases. This is also the case with the research themes of the Liggins Institute. Understanding the genetic mechanisms through which the early life environment determines an individual's adult body type and health profile is now becoming a computational problem, especially in the area of brain development.

The Malaghan Institute initiates cutting edge immunology research into the prevention and treatment of cancer, multiple sclerosis, asthma, arthritis and infectious diseases. The development of drugs to treat and prevent these diseases is determined by pharmacological studies. It is now becoming a daily exercise that large scale computers are being used to investigate how the human body transports and reacts to these sometimes-toxic drugs. However each person can react differently to toxicity so that combinations of drug parameters need to be investigated to ensure low risk. With parameters simulating the drug transport in the body numbering into the thousands the use of high performance computing technology is essential if we are to maintain a high quality of research and drug development within New Zealand.

---

[11] NZ Genomics has some of its HPC needs already funded, and is committed to utilising BeSTGRID for its Grid Computing needs, and its resources will be fully integrated with the HPC resources to be established with this proposal.

3.4.2 High-Value Manufacturing and Services

The success of developing new niche products within the nano-technology field relies on a tightly coupled programme of experiment and validated numerical models. Large-scale computational resources will enable New Zealand to maintain its international status in this growing and important field at the nanometre scale. In doing so it provides the basic science from which niche products are designed enabling economic growth. As described below these models require some of the most powerful high performance computing architectures yet designed.

Nanotechnology and materials is a rapidly developing and highly active field worldwide, across several traditional disciplines in the physical and biological sciences and engineering. Theory and simulation are essential components in understanding new science and designing new technologies in a situation where materials behave in radically new ways, and many important quantities are still inaccessible to direct experimental measurement.

Nanotechnology encompasses many outstanding scientific and technological challenges across a range of disciplines, including: (i) the development of new, smaller and faster information-processing devices that will be needed to extend the four decades of exponential reduction (in size and price) since the development of the transistor and the integrated circuit (“Moore’s Law”); (ii) the mechanistic understanding of the part that biological macromolecules and their assemblies play in the machinery of a living cell (“proteomics”); and (iii) the combination of these two fields to sense and monitor biological processes and inform medical diagnosis, on an unprecedented fine scale.

There are several types of nanostructure whose properties are of interest as the building blocks of this new science: (i) *quantum dots and nanocrystals,* (ii) *quantum wires and nanoscale logic devices* and (iii) *nanoscale electromechanical devices.* Such building blocks are increasingly being assembled into complex combinations; for example, nanoparticles coated with organic molecules are being proposed as bio-compatible markers, in photo-voltaic devices, or as electronic components. It is apparent that a range of techniques from different simulation fields (for example, computational chemistry, materials simulation and biomolecular sciences) is required to analyse this kind of system. All of the above—by virtue of the length scales and complexities involved in the model—require large-scale HPC. BlueFern currently supports a large part of this research due to the high processor count technology that is required but it is clear that simulations are constrained by the present environment.

The MacDiarmid Institute, a world leader in nano-materials, uses theory, experimentation, computational simulation and data-driven science to inform and stimulate experimental investigations into areas such as high-temperature superconductors and semiconducting and metallic nitrides, molecular magnets, solar energy and electroluminescent materials, critical to the success in a number of applications including solar cells, organic light emitting diodes (OLEDs) similar to those lights we see on bikes everyday, along with sensors and magnets. The Institute investigates the generation of smart, functional, viscoelastic matrices. Ubiquitous not only in biology, soft materials and complex fluids play a crucial role in industrial arenas as diverse as oil recovery, food technology, cosmetics and personal care products, electronic devices, and biotechnologies, such as microfluidics and targeted drug delivery. Fluid dynamics requires state-of-the-art computational resources to investigate the strange properties of fluids operating at the nanometre scale. It is now common practice to simulate these fluid properties using discrete methods (modelling how single molecules interact with each other to produce a continuous fluid) thus requiring large scale, high processor count HPC technology.

Improvements in computational resource are expected to be particularly crucial in the next five years since many of the methods of computational chemistry and nano materials are finding increased applicability in cognate application areas such as atmospheric science and biology, and these applications present extreme demands. An additional development is the integration of different (accuracy, length-scale, time-scale) methodologies to enable unified whole-system approaches to predicting and understanding the properties of matter, and such activity cannot be contemplated without an ongoing commitment to quality HPC facilities. Lattice Quantum Chromodynamics (QCD), the quantum field theory of the strong interaction, is the pre-eminent particle physics application requiring huge computing capability; world-leading research is impossible without world-class HPC resources.

3.4.3 Hazards and Infrastructure

New Zealand's location astride a plate boundary in the roaring forties inflicts a unique mixture of weather and geological hazards. Internationally, one of the main uses for HPC is its application to weather forecast models, and this is a prime application for New Zealand with its wide range of weather dependent industries. The weather models however are just the heart of the hazards forecasting chain, and by coupling them to catchment, river flow and then inundation models, flood forecasting can be achieved.

Likewise the weather model output will be used to drive ocean wave and storm surge models to develop integrated forecasts of our coastal hazards. In the quiet (non-hazard) times the same models can be applied to water resources problems, and to downstream production forecasts for our primary sector industries. Other applications include forecasting air quality and the risk from airborne hazards (e.g. foot and mouth disease, chemical contaminants), and modelling the propagation and inundation due to local and distant- source tsunami.

New Zealand's unique geological setting brings with it the challenge to understand, model and predict the impacts from earthquakes. The challenges for realistically modelling seismic events involve working at finer spatial and temporal scales, and deeper into the subsurface. Computational simulations must address seismic wave propagation from extended, complex faults, in three-dimensional irregular media. Lack of a suitably powerful HPC platform greatly restricts the fidelity of current efforts, and thus also their predictive power in assessing likely impacts on population and infrastructure. Building on the fine-scale simulations and scientific visualisation produced at the Southern California Earthquake Center[12] will allow a deeper understanding of seismic event outcomes, and thus aid in disaster planning.

eScience has an important role to play in combining our hazards research and data to provide an "*all hazards*" approach to evaluating risk and hence making decisions on the level of investment in mitigation. This approach is being followed in the recently formed Natural Hazards Research Platform, where the properties of the various perils, their interaction with communities and infrastructure, the resulting damage and mitigation options can all be explored in an integrated manner, creating tools for a wide range of end users.

3.4.4 Environment, Climate and Economics

New Zealand's terrestrial and marine ecosystems are extremely vulnerable to invasive species, which can cause significant economic losses if unchecked. Understanding the invasion process and the impacts of planned interventions requires highly-parallel simulations at fine spatial and temporal scales. Current efforts are stifled by the time such models take to run on available platforms. In order to model entire regions at a fine enough scale to remove the massive errors introduced by generalising the data, at least an order of magnitude increase in computing power is required. Invasive species modelling is just one example of many simulation challenges involving interactions between NZ's natural ecosystem and human society. The fates of these two systems are intimately entwined and the HPC offers one of the few viable approaches to resolving the complex suite of issues we face.

---

[12] See http://www.scec.org/ and visualisations at http://www.scec.org/resources/movies.html

Future modelling and prediction of environmental change will rely on sophisticated linked models of environmental components, and hence on HPC—to drive physical models such as climate and ocean models, and on eScience—to combine environmental datasets and link the component models such as ecosystem models.

As global economies continue to grow, so their impact on the environment increases and the environment's ability to provide the economic services on which humanity depends is put to test. Climate change is just one aspect of the constant interplay between humanity, economic growth and the environment. Climate is a globally shared issue that can be addressed by global modelling efforts in which we play a part. Our individual and collective impact on the terrestrial environment is also a global issue—our lifestyles are dependent on goods and services transacted globally with globally dispersed highly diverse local terrestrial impacts. Modelling and understanding these interactions is going to be fundamental to the future peace of the planet. In New Zealand we are just starting to assemble the capability to build the datasets and models to address this challenge; this endeavour will rapidly grow to a point where access to the scale of HPC facilities proposed here will become essential.

Climate modelling globally has relied on HPC to provide the resources necessary to perform long duration (decades to centuries of model time) runs of climate models, such as those used in IPCC assessments. Similar models will be applied in New Zealand to "downscale" the global model output to provide regional climate forecasting, with a focus on seasonal predictions and climate extremes, including the effects of extremes on river systems and the coast. HPC will provide the capability to run state-of-the-art models with a resolution suitable to capture regional variations for application across New Zealand's climate sensitive primary production sector.

In a similar manner, ocean models will be used to understand and model ocean variability around New Zealand and in future its impact on our ocean and coastal ecosystems, their productivity, and the likelihood of events such as algal blooms. This modelling will also strengthen our relationships with Australian researchers who are using the same climate models and with whom we share an atmospheric and oceanic "boundary".

Current generation climate models prescribe the atmospheric composition (e.g. greenhouse gas concentrations). The climate however affects the chemistry and hence concentrations, which in turn affect the radiation (energy) and hence the dynamics. Chemistry-Climate models that include atmospheric chemistry explicitly have been developed recently and will be applied to increase the fidelity of future climate modelling.

Exciting opportunities exist for combining related research areas. For example, coupling of climate (and weather) models to crop and production models in the primary sector can be used to predict production, harvest, storage and transport requirements, and linking diverse information on urban form, transport and domestic energy usage, and air quality will be used to understand the relationship between these previously disparate fields and hence inform future urban planning.

3.4.5 Energy and Minerals

New Zealand's needs are to have the information and tools to understand our resources (potential and realisable), to deploy them while minimising impacts, and to operate them efficiently. With a target of 90% renewable electricity by 2025, there is a focus on evaluating the resources, often in the absence of comprehensive instrumentation.

HPC based models are used as a surveying device to help select sites for wind-farms. Wind turbines require steady streams of air, as too much turbulence can damage their blades or stop them turning. Unfortunately many sites are complex, mountainous and prone to turbulence. By the conventional methods, putting up weather stations and collecting data, finding a suitable site can take a year. Using HPC modelling techniques combined with field observations reduces that evaluation to a single month, lowering costs of site identification in advance of local testing and allowing large numbers of designs to be evaluated *in-silico*.

In a similar manner other renewable sources such as marine energy will be evaluated using high-resolution wave and tidal current models, both dependent on HPC to obtain information at sufficiently small scales to reveal both the mean flows and possible damaging turbulence.

Advanced geothermal prospecting techniques and the modelling of geothermal reservoirs leads to greater understanding of the resource, extraction rates and sustainability. This is again a complex 3-dimensional problem that requires significant computational resources and sophisticated analysis and visualisation tools, along with services to integrate massive quantities of data. eScience applications can be used to include a wider range of information, for example mapping the existing (and required) infrastructure and its relationship to the range of distributed resources, for both electricity and our on-shore and off-shore mineral resources.

The requirement to accommodate a wide range of renewable energy sources places strong operational constraints on the New Zealand electricity system. A key element to enable robust operation is the ability to forecast the supply, demand and transmission conditions (capacity is dependent on temperature, solar radiation and wind). The linked forecast models described in Hazards and Infrastructure can be augmented to provide the relevant information. In particular, fine scale wind forecasts will in future be provided by high resolution weather models, enhanced by embedded very high resolution fluid flow models operating at the scale of individual wind turbines. Producing timely results at these small scales cannot be achieved without matching HPC resources.

3.4.6 Health and Society

On the surface there would seem to be little connection between the use of HPC in the research environment and the nation's health. Yet in fact over the past ten years HPC has become not only a tool which can help clinicians but a necessary infrastructure: e.g. it supports patient specific tools for investigating the best possible coronary or peripheral bypass; imaging protocols that provide previously unseen anatomical pathologies; and simulations of the human vasculature enabling non-invasive diagnoses of vascular diseases. The list grows exponentially. Over the next few years we will see more patient-specific tools being generated and this requires considerable support from the HPC research community. As a developed nation our health burden is a significant portion of government spend. HPC has the ability to help provide a richer environment for clinical work at an affordable cost, with potential for greater cost-efficiencies in treatment.

Water quantity and quality are crucial to our continued primary production and the health of the nation. Comprehensive modelling and prediction of our water supplies relies on running models for weather and climate (the renewable water supply), catchments and flow, which must then be linked to ground water models—all require HPC. As does the study of the leaching of nitrogen, viruses and bacteria from septic tanks into ground water supplies. These problems also require eScience infrastructure to maintain collaboration across a wide range of research groups, institutions and stakeholders.

For New Zealand within the health and society area there are a number of research communities with significant computational requirements, namely genomics/bioinformatics, biomolecular and physiological bioengineering simulations, although these fields overlap significantly in many places.

3.4.7 The proposed Square Kilometre Array

The Prime Minister of New Zealand indicated in 2009 the willingness of New Zealand to support Australia in a bid for the international Square Kilometre Array radio-telescope facility. This project is not on the same timeline. However, regardless of its location, New Zealand needs high performance computing capabilities and capacity to be able to play a role in the processing and use the data associated with this kind of landmark international facility.

While this proposal does not have the budget to support the SKA, having an existing infrastructure through which such investments can be realised would streamline the implementation of future infrastructure needs for the SKA.

3.5 Summary of Wider Science Community Benefits

In addition to the HPC capacity provided by this investment, there will also be many tools, services and applications developed that will offer significant benefits to a large number of research communities and institutions, improving efficiency and reducing some of the technical barriers to collaboration between institutions. These services will be made widely—and freely—available; most can be provided on demand via the grid middleware at no cost. Examples include:

Science benefits that will accrue to the NZ research community as a whole are:

- All institutions will have access to a powerful, integrated, scalable computational ecosystem and opportunity to gain benefit.
- Coordinated provision of world-class HPC and eScience infrastructure for nationally strategic research communities, including current and projected large-scale, national and international science and technology projects, such as the grand challenges of weather and climate prediction and the Global Research Alliance.
- Other research communities will have opportunities to benefit also, such as archaeology—with its massive data collections, and economics—with its computationally intensive financial risk models.
- Efficiencies in research funding via shared access to expensive HPC equipment, shared research infrastructure and associated productivity tools, data collections and experiments. Continued piecemeal funding will cost the nation more in the medium and long term and lacks the opportunity to develop the capabilities described.
- Smarter interfaces between NZ-based researchers and their international counterparts. Access points to trans-national European, Australian and US HPC facilities ensuring the NZ research community is connected with the world’s most advanced HPC communities.
- NZ able to play its part in the international eScience communities by making our data, applications and results available to the emerging global research infrastructure, to address important issues such as climate change impacts, plate tectonics, and bio-security.

Science benefits that will affect individual researchers are:

- ‘On the job’ training of our young scientists, engineers and IT personnel in the use and support of HPC, grid middleware and eScience-supported research.
- Scientists sharing *their equipment* (e.g. HPC resources), *their experiments* (workflow environments, modelling languages), *their data* (data repositories, web and grid data services), *and their understanding* (virtual collaboratories, knowledge capture tools), thus reducing research costs.
- Rapid, customisable assembly of services, applications and workflows thereby reducing costs of developing applications, supporting infrastructure and computing needs. Increase the ‘science as a service’ approach to sharing research outcomes with relevant organisations.
- Easy job submission tools that run simply from the user’s desktop, so that researchers are not forced to spend time trying to get to grips with large facilities and different software.
- Collaboration tools such as Wikis and Virtual Organisations, which once set up can be provided to any community that wants to make use of them.

- Storage infrastructure and associated tools to manage, organise and serve large data collections that are important science assets, such as nationally significant databases.
- Access and Identity Management, to ease the problems of accessing remote equipment and help to encourage cross-institutional collaboration.

4 THE ECONOMIC CASE

To help drive this economy, NZ researchers need to keep pace (and collaborate) with overseas colleagues whose needs for research capabilities are often better provisioned. A healthy digital economy requires an innovative infrastructure for access to computing resources, data collections and tools to support scientific discovery, information analysis, simulation, visualisation, and sharing of results, in support of our research communities[13].

However, our current HPC and eScience infrastructure is ad-hoc, under-funded, over-utilised and can offer only short-term service guarantees to its user communities. To date, the lack of a national plan, combined with financial and research funding obstacles have impeded wide science sector access to capability-class computing in New Zealand which, in turn, has discouraged its development. As a consequence, many scientists have tended to only tackle problems that are manageable on desktop systems—leaving a growing void in New Zealand science and technology capability. This is creating an opportunity-cost to New Zealand in terms of lack of advancement of a diversified economy, slower product discovery and development, and loss of innovative capability overseas (both top talent and businesses).

Large-scale and centrally-coordinated eScience infrastructure is now being implemented and used throughout most science-led economies, including all of Europe, the USA, Canada, Australia, Japan, Korea and even parts of India. Appendix 1 provides some comparisons of several international investments. The pan-EU initiative describes eScience Infrastructure as a:

> *"…new research environment in which all researchers—whether working in the context of their home institutions or in national or multinational scientific initiatives—have shared access to unique or distributed scientific facilities (including data, instruments, computing and communications), regardless of their type and location in the world."*[14]

This section describes the economic benefits of creating a stable, well-equipped HPC and eScience platform in New Zealand, and provides strong international, independent evidence of the advantages that follow. It is organised into five themes:

1. The *first theme* (4.1) explains why New Zealand should improve its own national HPC capabilities, rather that purchasing them from offshore.
2. The *second theme* (4.2) describes the current situation in New Zealand, with some successes and great potential for success, but constrained by a fast-growing need for more advanced science infrastructure.
3. The *third theme* (4.3) summarises the savings and other benefits that can be achieved by a collaborative, sector-wide approach, rather than each institution struggling to provide its own increasingly sophisticated science infrastructure.
4. The *fourth theme* (4.4) shows the economic benefits for an offshore economy (Queensland's) that were achieved by investing in coordinated provision of eScience.
5. The *final theme* (4.5) describes the results of a comprehensive, independent survey conducted in the UK to assess the value of HPC and eScience infrastructure to the research community, and highlights the productivity increases and efficiency gains accrued.

---

[13] Paul Jeffreys, eScience Centre, U. Oxford.

[14] http://cordis.europa.eu/fp7/ict/e-infrastructure/home_en.html

4.1 Analysis of the International HPC Market

Buying 'HPC On Demand' or HPC in the Cloud' from international sites is now possible, and we have investigated this option. There are several reasons why these systems would not be beneficial as an alternative to a fully integrated NZ-based infrastructure, and we set these out below. We note also that the *NZ MetService* recently reached the same conclusions when evaluating their HPC options[15].

In carrying out this analysis a specific focus was given to the recently announced Amazon HPC-on-Demand resources "EC2" and "CCI", although other solutions and services were reviewed from SGI, IBM and other vendors similar offerings. The following discussion highlights the fundamental differences between the needs of the NZ research sector for HPC and those services currently being offered in the Cloud.

HPC in the Cloud is provided on commodity computers that lack the specialised networking fabric that knits Supercomputers and Clusters together internally to provide the low latency high performance characteristics required. In the middle of 2010 Amazon released Compute Cluster Instances (CCI), which addresses this need in a narrow way by providing faster networking between CCI nodes. These environments have been designed so that many lightweight and short-lived tasks can share a single computer without knowledge of the underlying hardware platform. This is the inverse of the need within the HPC community, where the need is for a single task to share the equivalent of many computers, intensively, and therefore highly optimised for the specific hardware platform being run on. The hardware platforms required by the NZ research sector are varied, whereas vendor services like those offered by Amazon are usually offered on only one hardware platform. The resulting design differences have meant the Cloud style services aren't suited to HPC needs. The narrow focus of these systems won't provide the varied requirements for a NZ infrastructure. However, the flexibility of HPC provision envisaged for NeSI would allow for Cloud computing to play a significant role in the emerging infrastructure if it becomes more viable.

There is a cost consideration too, as these systems are priced to be attractive to the individual or small-scale user. This might be suitable to a lone researcher or small research group, who mostly won't have the funding to purchase HPC resources for their own use, and certainly not at the scales possible when large capital investments are coordinated across a nation. This coordination provides considerable purchasing power, and allows for the infrastructure to be established and delivered to researchers at a significantly lower cost, when compared on a per computer processor cost per hour.

Additionally HPC On Demand has obvious issues with the time taken to pass data between its compute nodes[16] and the data storage disks that are sometimes remote from the actual computer itself. There are costs associated with hosting data or retrieving data from the infrastructure and questions over the security of the data housed. With the geographic remoteness of NZ and the small levels of international bandwidth on our research network, these costs are further compounded by performance issues. The complexities of the EC2 file system makes the use of Compute Cluster Instances not trivial and any user adopting the HPC in the Cloud will need to contract expert independent services (an additional cost). Hence the overseas capability would not be able to provide an integrated service to users of HPC.

At this point in time only small-scale cluster architectures are available from HPC on Demand providers, hence no high processor count or more advanced computing architectures. The one exception is an IBM BlueGene available as a resource from their centre at Rochester (USA); however this is prohibitively expensive. Furthermore, there is a reduced amount of software available for users on such commercial platforms, compared to that already available for NZ researchers.

---

[15] MetService rejects cloud: http://computerworld.co.nz/news.nsf/technology/metservice-rejects-cloud

[16] (one instance = a pair of quad-core Intel "Nehalem" X5570 processors with a total of 33.5 ECU (EC2 Compute Units), 23 GB of RAM, and 1690 GB of local instance storage)

Australia has several compute resources, for example the new Victoria Life Sciences Computation Initiative (VLSCI) at University of Melbourne. However Australia could not provide enough capacity for NZ usage and NZ researchers would be at 'the mercy' of both the trans-Tasman broadband KAREN link and the prioritisation of the Australian research community over NZ researchers. Indeed the data generated by NZ researchers would quickly saturate the single pipeline across the Tasman. This would not be the case when using the various pathways for KAREN within New Zealand nor when providing our own capacity.

Survey evidence suggests that **local** services, including support services, are integral to people taking up the opportunity to improve their research by use of HPC. People need to trust the service, and this is much easier with the local-support service model proposed in this infrastructure model.

More importantly, without the HPC infrastructure in NZ, we will not develop the skills to operate and leverage these platforms effectively, experiment upon them widely and integrate them tightly with the needs of our own researchers.

Benefits of overseas linkages with an increased NZ capability
In fact the opportunities for leveraging offshore investments in HPC actually increase as we build additional capability and skills locally, as these skills equip our researchers to make optimal use of local and overseas HPC, such as the new Australian platforms in Melbourne (Victoria Life Sciences Consortium) and Perth (IVEC). Creating connections within the international HPC community has a double benefit. It provides dissemination of best practice for HPC services to the researchers and allows entry into the use of massive HPC infrastructure far larger than most developed countries could afford—though not before we have proved that we can scale up our research on the locally-available HPC platforms. We give below two examples that support the above statements.

HPCWorld is a consortium of five Supercomputing Centers (CINECA, BSC, FZJ, SDSC and BlueFern) plus GENCI, the French national agency in charge of HPC coordination[17]. The aim of HPCWorld is to survey and analyse the models used worldwide by the HPC centres and initiatives to manage the requests for allocation of resources coming from researchers and their communities. The results are then presented and offered as a best practice to all HPC initiatives and centres, as well as to scientific communities that may take advantage from it. Since BlueFern is already a member, NeSI will gain immediately from this collaborative grouping.

The U.S. Department of Energy (DOE) Office of Science provides a portfolio of national high-performance computing facilities housing some of the world's most advanced supercomputers. These leadership computing facilities enable world-class research for significant advances in science.

Open to researchers from academia, government labs, and industry worldwide, the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program[18] is the major means by which the scientific community gains access to some of these supercomputers. The program aims to accelerate scientific discoveries and technological innovations by awarding, on a competitive basis, time on supercomputers to researchers with large-scale, computationally intensive projects that address "grand challenges" in science and engineering. Powerful, leadership-class computing systems at DOE's Argonne and Oak Ridge National Laboratories are providing 1.6 billion processor hours to 69 INCITE projects around the world for access in 2011. These allocations support high-impact advancements by providing computer time and resources for one to three years.

This is of particular importance to researchers in New Zealand accessing resources through NeSI since they can utilise NeSI resources in order to produce sufficient evidence to the INCITE programme of ability to scale up the problem to their advanced architectures and processor counts.

---

[17] www.hpcworld.eu

[18] http://www.doeleadershipcomputing.org/

Since the computational infrastructure at Argonne is of the order of 100,000 of processors, it is crucial to show that codes can scale to this number of cores. New Zealand thus has an opportunity to access free, massively-scalable HPC if we have the skills to use it.

4.2 Benefits Specific to the Government's Economic Growth Agenda

HPC capacity and eScience services will be targeted towards the needs of New Zealand's most strategic science communities. These communities have been chosen because of their importance in contributing to an innovation-led economy. Where there is an opportunity to prioritise access to resources, the **potential economic benefit** of the research proposed will play a role in allocating those HPC resources.

As Figure 4 (in the previous section) shows, these communities align well with New Zealand's RS&T priorities, and **will help support a diversification of the NZ economy** by facilitating research and development across all sectors supporting the Economic Growth Agenda, from agriculture to health, materials, biological industries, energy, environment, climate and hazards mitigation. The Science Case, in the previous section, describes some of these benefits in terms of the research needed to achieve new breakthroughs that in turn will stimulate economic growth.

More detailed explanations to illustrate how the infrastructure will facilitate research outcomes are provided in selected examples below.

The 'opportunity pipeline model' (Figure 5) shows the potential for future growth if good science is able to progress to maturity. At the moment (shown by the 'current state' arrow), this funnel is narrow: lack of access to appropriate research infrastructure artificially restricts the flow of projects to market maturity. Taking a national perspective, we have patches of demonstrated capability—see below—but not enough capacity. The proposed eScience infrastructure will allow more projects to progress along the development pipeline, and at an increased pace, since they will leverage a greater depth of eScience services and computational platforms. External benefits include economic gains, public good benefits and stronger international collaborations. Example projects at each stage in the opportunity pipeline are discussed further below.

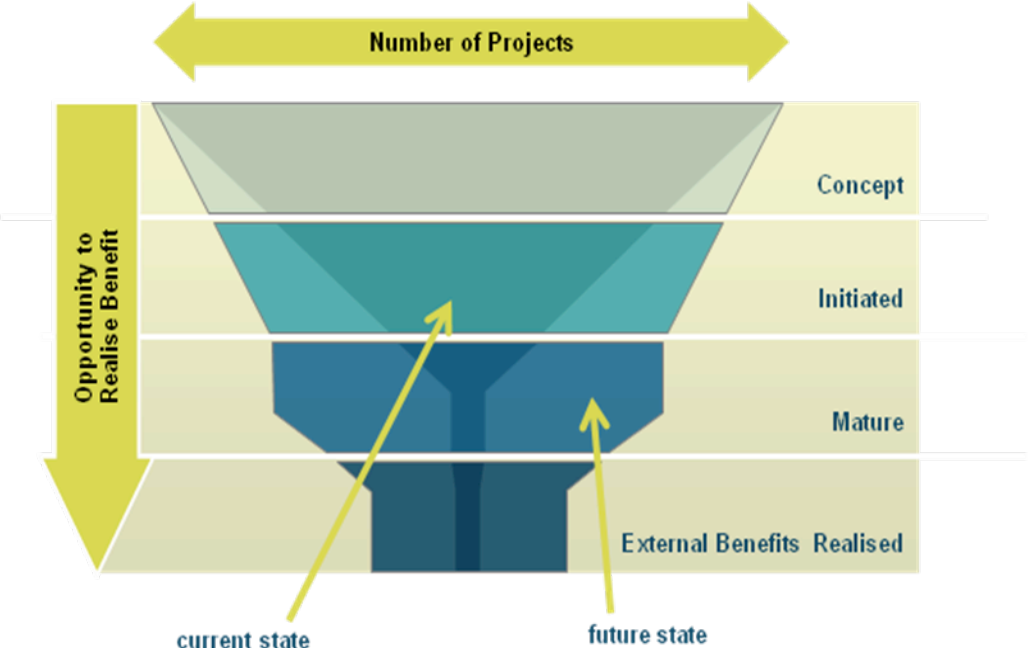

Figure 5: **Opportunity Pipeline Model**, showing the path from concept to market for science-related applications.

Examples of initiated and mature eScience applications are:

1. **Initiated—Precision Drugs for Cancer**: Using eScience techniques, Dr Jack Flanagan at the Auckland Cancer Society Research Centre has developed a method with the potential to scan 50,000-100,000 ‘virtual’ drug candidates within one day, far outperforming any ‘wet’ lab experimentation. This will speed up the economic benefit to be gained by drug discovery.
2. **Mature—Better diagnosis and treatment of injury and disease**: The Auckland Bioengineering Institute (ABI) led by Professor Peter Hunter. ABI is a mature, large-scale research facility with demanding eScience requirements; global collaborations and international investment. It has multiple international funding contracts, and is developing novel devices and models, that will create savings and better quality spend in the health sector overall.

Additionally, many of the tools, services and applications developed as part of the infrastructure will offer significant benefits to individuals, research communities and institutions, providing access to HPC, improving efficiency and reducing some of the technical barriers to collaboration between institutions.

The examples below highlight the economic impacts of improved research capability across a number of fields.

4.2.1 Biological Industries

It is critical to New Zealand’s primary industry that we remain at the forefront of DNA technology-based farming and agricultural research and practices. As our agricultural trade competitors are becoming more sophisticated, so we need to retain a leading edge in our primary productive sectors. Advanced, high through-put genomic, proteomic and metabolomic analyses enable a more holistic understanding of biological processes, which are highly data intensive, hence requiring HPC. This is necessary to enable advanced productivity in the primary productive sectors.

There are huge commercial incentives to scientific success across the life sciences. Examples include the ‘Primary Growth Partnerships’ for agriculture, aquaculture, horticulture and forestry. Examples are, where eScience (i) creates significant new ways to optimise the value of milk production by modelling the cow rumen (the “engine room of the NZ economy”), (ii) enables linking the bovine genetic (haplotype) factors and environmental factors (including diet) with quantitative models of phenotype that include fatty acid production and milk production and (iii) linking environmental and production research so that weather and climate models can be applied to predict pasture growth and ocean modelling can predict the availability of nutrients for aquaculture.

One example that depended on genomics/bioinformatics is a 2007 joint venture between the University of Otago and Ovita, a biotechnology company jointly owned by Beef and Lamb New Zealand and AgResearch. The overall purpose of this venture was to advance the identification of the gene variants underlying meat and wool production traits, as well as those that cause susceptibility to disease in the sheep genome. Researchers can use these variants to develop tests for breeders to identify and select animals that have superior meat production and wool quality, as well as a resistance to disease, hence increasing the economic outputs of these producers.

While this is traditionally an area of research strength in New Zealand, new computational facilities are needed to provide a step-up in the ability to construct genomes, share and annotate data pertaining to functions of genes, and discover applications of the knowledge. Without the computational facilities and bioinformatics analysis tools, the competitiveness of this area of strength for New Zealand will decrease, compared to overseas where greater computational and eScience capabilities are available.

4.2.2 High Value Manufacturing and Services

High performance computing underpins research into advanced materials and nanotechnologies, which will be key contributors to New Zealand's high value manufacturing sector, currently worth more than NZ$5bn in exports per annum. For example, the MacDiarmid Institute's research has led to a number of spin-out companies in high tech manufacturing including Magritek, HTS-110, and Anzode. The institute also works with small start-ups such as Izon, larger established manufacturers such as Fisher and Paykel and Fonterra, and multinationals such as IBM.

High performance computing is used to provide a virtual laboratory where new materials can be designed, and virtual experiments can be conducted. Such virtual experiments often cannot be performed at all in real laboratories, (e.g. because they deal with nano-structures), so they can provide crucial information that could not be obtained otherwise. They are also vital for interpreting the results of real experiments, allowing researchers to test models for materials and devices and extend these models to novel materials. Novel materials elicit high prices when used in high technology manufacturing (such as electronic components and devices) and will underpin the transformation of New Zealand's manufacturing sector. Finally, high performance computing is essential for device design, allowing rapid prototyping and device optimisation, achieved through modelling novel material characteristics and their interactions.

High performance computing resources are necessary for New Zealand researchers to stay at the cutting edge of fundamental and applied materials sciences, providing graduates with training in the best modern scientific techniques.

4.2.3 Hazards and Infrastructure

New Zealand's unique geological setting (i.e. on the boundary of two tectonic plates) means that understanding the impacts of earthquakes is vital. To obtain meaningful data, computer simulations must address seismic wave propagation from extended, complex faults, in three-dimensional irregular media. Lack of a suitably powerful HPC platform greatly restricts current efforts, reducing their predictive power in assessing likely impacts on population and infrastructure. Creation of an integrated hazard forecast system that depends on HPC, such as NIWA's EcoConnect, would have immediate benefits in terms of anticipation of major events.

Coordination with the Natural Hazards Platform, and drawing on the fine-scale simulations and scientific visualisation produced at the Southern California Earthquake Center[19] would allow a better understanding of seismic event outcomes, and aid New Zealand in planning for natural hazards and disaster relief. Visualisation and data collection management services are also required, along with infrastructure for more effective sharing of data, models and results.

Flooding is New Zealand's most frequent natural hazard, with average insurance payouts of $70m per annum, and economic (direct and indirect) losses much higher. Cost benefit analyses performed overseas (e.g. in the UK) have demonstrated the payback for investment in prediction systems, on a long time scale to inform planned mitigation, and for short-term emergency management. Much recent HPC use has focused on implementing and running state of the art models to improve the predictions of both quiescent and extreme events at a range of time scales for application in the relevant sectors.

4.2.4 Energy and Minerals

Renewable sources of energy are important both to New Zealand and the global community, and our mineral endowment has wealth creation potential. Both areas benefit from applications of HPC and eScience during exploration and operational phases. Exploration of our renewable energy resources is greatly enhanced through assessment of the wide range, variability and location of

---

[19] See http://www.scec.org/ and visualisations at http://www.scec.org/resources/movies.html

potential resources (e.g., wind, geothermal, and marine) using models (atmospheric, geothermal reservoir and coastal hydrodynamic models) to substitute for extensive and expensive measurement campaigns. Offsetting carbon-based energy requirements reduces the costs associated with New Zealand's obligations under the Kyoto Protocol.

Wind energy provides a good example as this research area requires significant simulation activity, using HPC models as a surveying device to help select sites for wind-farms. Wind turbines require steady streams of air: too much turbulence can damage their blades or stop them turning. Unfortunately many potential sites for wind farms in New Zealand are complex, mountainous and prone to turbulence. Finding a suitable site through conventional methods (i.e. putting up weather stations and collecting data) can take up to a year – a costly delay. Using HPC modelling techniques, combined with field observations, reduces the evaluation time to a month. This offers significant benefits, including lower costs, and enables a number of designs to be compared and evaluated *in silico*.

Forecasting most renewable energy sources relies on weather forecasting models, which are dependent on HPC. Application of these models enables better prediction of hydro and wind resources, supporting energy planning and also the operational integration of variable renewable sources into national (and local) networks to optimise the use and storage of the resources, while reducing risks of unmet power demands. This can reduce New Zealand's commitments under the Kyoto protocol, while setting the standard internationally for optimising the proportion of renewable energy used in a nation's electricity production.

Advanced analysis of seismic, topographic, bathymetric and magnetic data is required to understand the nature of under-sea and terrestrial resources and the economic potential for extraction. HPC, coupled with appropriate software, provides advanced analysis, enabling good quality decisions on block values and exploration, creating greater value out of leases or exploration permits. By working closely with the *GEON* and *AuScope,*[20] international eScience communities, a large number of geophysical exploration, modelling and tools for data integration and visualisation could be added to New Zealand's existing capacity.

4.2.5 Environment

The New Zealand economy is highly dependent on weather and climate, for primary production, renewable energy, and transport, now and into the future. Much of our recent HPC use has focused on implementing and running state of the art models to improve climate predictions at a range of time scales for application in the agricultural sectors. Prediction of water availability and drought forecasting is critical to the success of an agricultural economy. Accurate, long-term predictions of climate allow remediation measures to be put in place early if a drought appears likely.

Global scale climate models provide scenarios of the future climate out to decades, but usually have very poor resolution for New Zealand, e.g. 10 data points to encompass the entire country. Of the ways to downscale this information to appropriate scales for NZ climate (to distinguish the South Island west and east coast climates), a Regional Climate Model (the same climate model with much higher resolution, and hence greater computational requirements, nested within the global model) provides a physically consistent tool. The RCM model output is only the first step in generating benefit. Uses to which the RCM output can been put include: (i) informing Guidance Notes for Regional Councils on climate change, published by MfE, (ii) providing future Flood Guidance to MfE after linking the RCM to flow modelling, (iii) providing input to work by Scion on future forest fire risk, (iv) informing climate impacts and adaptation for agriculture (work by NIWA, AgResearch, Plant

---

[20] GEON and AuScope are geoscientific eScience infrastructures in the USA and Australia, respectively. They offer a rich set of geophysical analysis services that could be brought into New Zealand infrastructure to increase our capability to analyse and visualise subsurface structure. See www.geogrid.org and www.auscope.org.au/

and Food, and Scion), and (v) providing input into international applications on Pacific and Asian climate change.

This wide range of applications for end users is indicative of the value of these HPC generated model outputs.

4.2.6 Health and Society

Water quality (its uptake, cleanliness and distribution) is crucial to our continued reliance on freshwater resources – not only for agriculture, but also for the health of the nation. Although water quality is primarily an environmental concern, it is also important to other RS&T priority areas including the biological industries (primary industries and food), and health and society. For example, tools are urgently needed to model the leaching of pollutants into ground water supplies. Such modelling requires HPC access, along with data services to support analysis and collaboration across a wide range of research disciplines. The resulting economic impacts include a reduction in health interventions, productivity through improved health and environment, and better informed decision-making.

The applied impacts of biomolecular simulation are in the area of drug discovery. While we have a history of performing well in drug discovery, the hit rates are low and hence high risk. To reduce costs of discovery (while also broadening the molecular space) searched requires either considerable investment in wet labs for high throughput screening or developing virtual high throughput screening using HPC. This second approach is a much cheaper and more scalable approach, which can lead to shorter lead times for drug discovery.

4.3 Savings from Coordination Across the Research Sector

Investment in shared eScience and associated infrastructure produces an efficiency gain across the sector because it removes the need for each institution to provide the increasingly sophisticated research platforms that its scientists require. As the costs in equipping researchers for modern science continue to rise, national coordination offers increasingly significant economic gains; not acting so will cost more in the medium and long term.

As well as reducing costs, and encouraging collaboration, this approach also helps reduce the risks associated with small, stand-alone capabilities, and over-reliance on the single individuals who often operate them. Improved outcomes in research productivity and efficiency have already been achieved within BeSTGRID-linked institutions by sharing infrastructure and developing a connected workforce to support this infrastructure.

A national eScience infrastructure will therefore:

- Avoid unnecessary duplication associated with investments by individual institutions
- Develop a specialist eScience workforce, coordinated throughout New Zealand
- Facilitate country-wide research collaborations in RS&T priority areas.
- Offer a stable, mature and reliable capability, reducing risk across the sector that stems from under-resourced facilities.

The following are examples of synergies that can be leveraged with existing collaborative projects.

4.3.1 Examples of collaborations with existing research consortia

These examples provide a view of the wide range of research groupings that can be brought together to focus on key research and development questions for NZ, all with a critical dependence on modern computation, data management and communications to achieve outcomes beyond the capability of the individual partners.

The **MacDiarmid Institute for Advanced Materials and Nanotechnology** is a Centre of Research Excellence (CoRE) concerned with high quality research and research education in materials science and nanotechnology. It is formed as a partnership involving four universities (Massey, Canterbury, Otago and Victoria Universities), and two Crown Research Institutes (GNS Science and Industrial Research Ltd), with Victoria University as the host. Its open, collaborative structure allows researchers from other organisations to participate freely, with a significant cohort of MacDiarmid Institute researchers based at the University of Auckland for example. Since being established in 2002 a significant part of the MacDiarmid Institute research has involved analytical and computational modelling studies of new materials and systems using high-performance computing; the Institute has its own computing cluster and since 2008 there has also been strong utilisation of the *BlueFern* system at the University of Canterbury. This has been an outstanding success for the Institute's research programmes, and there is strong support for additional, national HPC capabilities to further advance Institute's computational materials science research.

The **Natural Hazards Research Platform** is a partnership between GNS Science, NIWA, Auckland, Canterbury and Massey Universities and private sector partner OPUS International. This partnership will integrate New Zealand's research into natural hazards, aligning funding sources to provide a coherent approach to the wide range of hazards we face, enhancing resilience, reducing economic losses and speeding recovery following severe events. The research and applications have a large requirement for High Performance Computing, with particular application in NIWA's weather driven hazards research and GNS Science's seismic modelling. The integration of natural hazard data in a similar manner to the Australian IMOS (Integrated Marine Observation System) would enhance both the research of the Platform and downstream end use by making all hazard data available in a unified framework.

New Zealand Genomics Limited is a collaboration between Otago, AgResearch, Auckland and Massey and has been granted substantial funding to establish an infrastructure to support genomics research. Key elements of this infrastructure investment include; high thru-put instrumentation, human resource capital and substantial computing resources. Over the initial five years of this funding substantial funds will be invested in high performance computing to support the genomics service. This investment will include; high level processing, commodity processing and substantial online and near online storage. This resource will be predominantly directed to the genomics community, and it is expected that it will substantially meet the majority of production related needs of this community. The board of NZGL considers it desirable to partner with this initiative, as NZGL will be a significant provider and consumer of HPC and of research-related computational services and middleware within the NZ research sector.

The recently-formed **New Zealand Agricultural Greenhouse Gas Research Centre** (NZAGGRC) was set up to address the significant challenge to reduce its greenhouse gas emissions from agriculture, with the benefits of reduced emissions costs, increased productivity and preservation of our green marketing advantage. The centre is a partnership with research and industry players AgResearch, Landcare, NIWA, Scion, Plant and Food Research, Lincoln and Massey universities, DairyNZ and Pastoral Greenhouse Gas Research Consortium. In a similar manner to the Hazards Platform, this partnership will integrate New Zealand's research on agricultural emissions, accelerating the research agenda, and will also serve as the local hub for the government supported Global Research Alliance. It will have substantial HPC and eScience needs related to modelling emissions and related strategy and policy research.

The **Virtual Institute of Statistical Genetics** was formed in late 2008 with the aim is to establish NZ as a world leader in advanced statistical methods for gene mapping by 2014. The Institute is a collaboration between Scion, Auckland, AgResearch, Otago, Plant & Food, ViaLactia, and ESR. The institutes goal will be achieved by combining New Zealand's top gene-mapping statisticians with geneticists to develop advanced methods for experimental design and analysis using pan-NZ computational platforms.

With the right eScience infrastructure in place, New Zealand researchers will leverage existing, global eScience networks, and participate in international collaborations more fully. eScience infrastructure has proved to be a major driver of improved efficiency and effectiveness within the research sector throughout the world. The following two sub-sections give detailed evidence of the successes achieved in Queensland and in Europe.

4.4 Economic Impact of eScience: the Queensland Case Study

A Queensland-based eScience infrastructure organisation, the Queensland Cyber Infrastructure Foundation (QCIF), in 2006 commissioned an independent economic impact assessment of their performance, operations and contributions to the Queensland economy.[21]

This report indicates that a potential five-fold return on investment over a 20-year period may be realised by funding eScience infrastructure.

QCIF is a consortium of Queensland universities formed to increase the State's innovative capacity through deployment and exploitation of high performance computing (HPC) and communications infrastructure. That infrastructure includes supercomputers, high-capacity data archives, visualisation, and networking capability.

The economic impact data from QCIF is extremely useful, as it provides insight into the potential economic benefits derived from New Zealand's investment in eScience infrastructure.

The Allen Consulting Group report identified impacts of QCIF in the following categories:

1. Economic benefits arising from:
   a. attraction of additional grants to Queensland,
   b. additional investment in Queensland,
   c. industry productivity gains and
   d. increased sales of related products.
2. Research productivity and related impacts.

To accommodate the scope of QCIF outcomes, the consultants used Economic Impact Modelling of returns from investment in QCIF, using Monash University's Centre of Policy Studies MMRF-Green model (Described in Appendix 3) to estimate QCIF's impact upon macro-economic indicators, such as gross state product (GSP), real consumption, employment, and attraction of additional industrial funding to Queensland. They also used assessments of specific research benefits, for example, traditional indicators of academic performance such as numbers of publications, patents and students, in combination with case studies of selected QCIF-sponsored projects. The report's findings are summarised below.

Return on investment

Across the period 2000-01 to 2019-20, the modelled net return from public sector investment in innovation made possible by QCIF to Australia was 68% in discounted terms (using 2.4% real discount rate) and 88% in undiscounted terms.

Across the same period, the modelled net return to Queensland was approximately 500% in discounted terms and even greater in undiscounted terms. The authors stated that their findings,

---

[21] Assessment of the economic impact of the Queensland Cyber Infrastructure Foundation to Queensland, Allen Consulting Group, 12 Dec 2006.

based on case study analysis, were consistent with the literature on returns on public investment in innovation.

Attraction of additional grants

In addition to $23 million in grants from the Queensland government, QCIF attracted approximately $62 million in grants from outside Queensland. The report noted that many of the QCIF's research projects had been in operation for less than five years at the time the report was commissioned, and some for less than a year. This suggests that the assessed outcomes, particularly scientific and commercial ones, are likely to represent only a portion of the total outcomes the projects might generate over the long-term.

4.5 Improved Research Efficiencies and Effectiveness: the European Case

In Europe, eScience infrastructure has been a major driver of improved efficiency and effectiveness within the research sector. Most European countries are involved in both national initiatives and pan-EU science infrastructures. For example, The Cook Report from Feb-Mar 2010[22] details the huge efficiency gains accrued in the Netherlands as a result of their highly-successful, pan-institutional approach to research infrastructure. To justify continued investment, various measures of impact have been studied. The following results were obtained from the EU study eResearch2020[23] conducted by a consortium led by Oxford Internet Institute leader Professor Ralph Schroeder. It included questions to users of established eScience infrastructure, to ascertain its impact on their research quality and costs. There were 388 respondents. The composition of the survey population is shown in Figure 6.

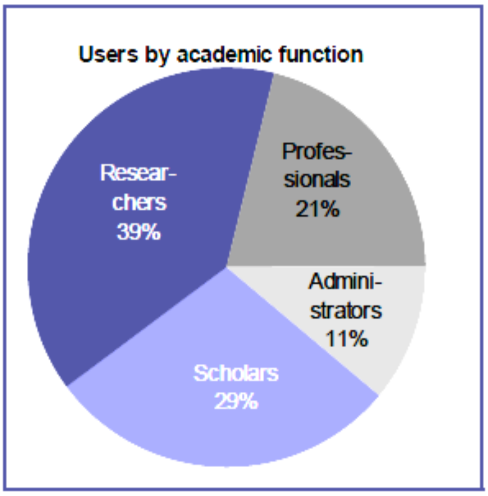

Figure 6: **Survey Respondents** by academic function.

Table 2 below shows the large, positive impact of e-infrastructure on the effectiveness of research across a number of different categories including reducing costs and increasing measurable outputs. Improvements ranged from 42% to 77%.

The EU study found that there was widespread agreement among users about the positive impact of e-infrastructures. For seven out of eight questions (shown in Table 2), more than 60% of respondents agreed that there was positive impact. The main benefits related to the speed of doing work:

- Accomplish tasks more quickly

---

[22] http://www.ictregie.nl/publicaties/nl_08-NROI-258_Advies_ICT_infrastructuur_vdef.pdf

[23] http://www.eresearch2020.eu/

- Access resources faster or better
- Work on previously intractable research problems
- Produce processes or analyse data faster and better

Almost as important was the ability to work on new problems that previously had not been able to be addressed due to technology limitations. Slightly less frequently, respondents agreed to positive effects on productivity, costs and quality. The lowest number of positive responses was on the acceptance of publications, perhaps due to the particular difficulties of assessing this impact.

Table 2: **eScience Infrastructure increases research effectiveness**. Results of an extensive survey of international research communities to quantify eScience impacts.

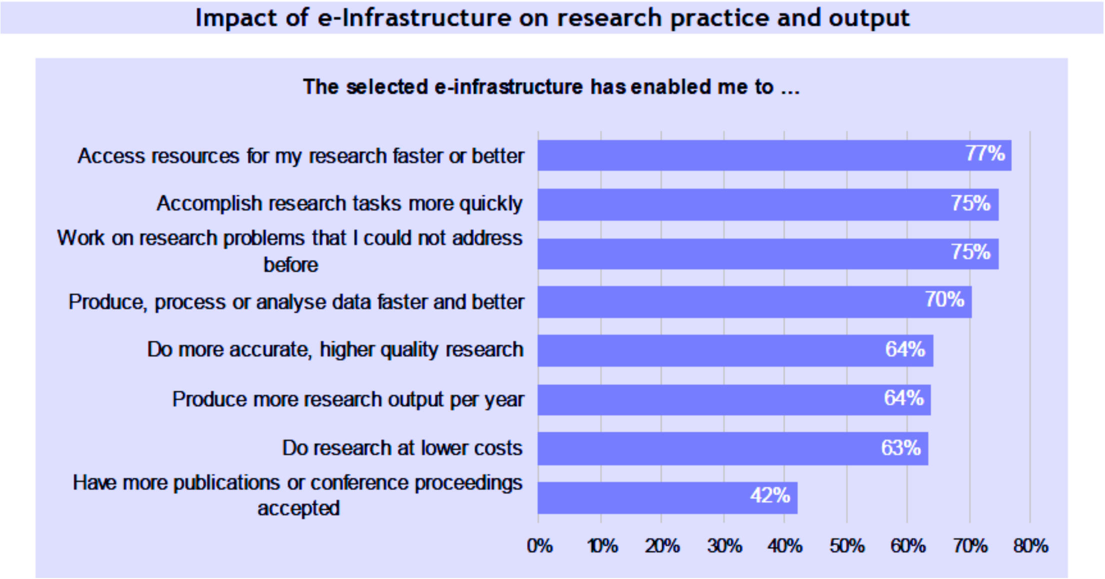

Table 3: **eScience increases collaboration** with both commercial firms and other researchers.

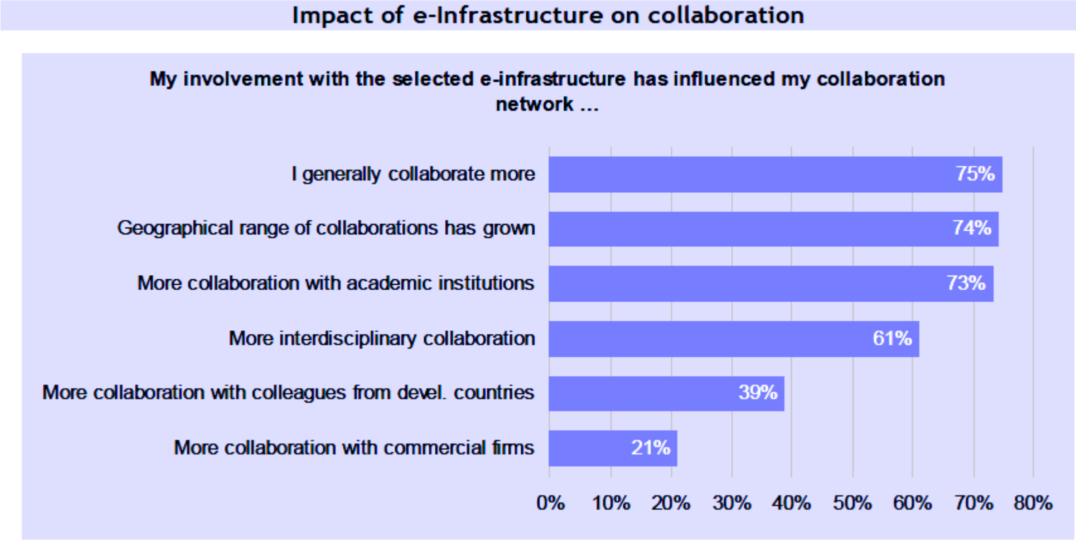

Table 3 shows the impact of eScience on collaboration. As well as its positive impact on research effectiveness, established eScience infrastructure facilitated richer and more useful collaborations. The results show that respondents reported a **21% increase of collaborations with commercial firms, and a 60% rise in collaboration with colleagues from other fields of science/work**, labelled ‘interdisciplinary collaboration’ in the table. Other benefits included increasing the geographical range of collaboration, enabling more interdisciplinary work and increasing with other countries.

5 BUSINESS AND GOVERNANCE

5.1 Legal Structure and Agreements

NeSI will be established via a legally binding agreement between the investing institutions for this purpose. A single "host" institution will be the legal entity through which the government provides funding via a Funding Contract written by the Ministry of Science and Innovation, and this host will provide legal, accounting, and other administration services to NeSI. Separate sub-contracts between the Host and other parties will coordinate capital investments between the four Principal Investor institutions, each delivering services under Service Level Agreements described within these sub-contracts, with capital grants made for establishment of essential capacities.

The responsibilities of the Principal investors will be defined within the following legal agreements:

1. **NeSI MoU**: Legally binding agreement between Principal Investors that incorporates the essential elements of this Investment case, including the **governance principles** (this section), the **Budget** (Section 6), the **Charter** (Appendix 5) and that defines joint responsibilities and objectives. This agreement effectively forms the collaborative venture referred to below as the 'NeSI partnership' or 'partnership', though these terms are not used in their legal sense.
2. **Funding Agreement**: Legal contract between MSI and the host institution.
3. Sub-contracts in the form of **Service Level Agreements** will be established between the host (as the contracting agent for NeSI) and each Principal Investor, in line with this investment case.

The role of the Ministry of Science and Innovation (MSI)

MoRST's successor, MSI, will monitor the Funding Agreement with the host and the performance of the co-investing partners; evaluation criteria are provided below in Section 5.5. MSI will be able to access information about sub-contract status, details, and performance. There will be a 'reach-through' clause to enable MSI to take action in exceptional circumstances, for example where a breach of sub-contract cannot be remedied within the partnership. An observer from MSI will be invited to attend all NeSI Board meetings.

The remainder of this section outlines NeSI's business and governance structures, processes, roles and responsibilities, all of which will inform the development of the above agreements.

5.1.1 Business and Governance Structures Evaluated

Several options for the legal structure and governance of NeSI were carefully evaluated during the formation of this investment case, including a review of several offshore and NZ-based infrastructures. Offshore investigations included: (i) *The Netherlands National ICT Research Infrastructure*[24] (ii) *Victorian Partnership for Advanced Computing* [25], (iii) *IT Center for Science (CSC) in Finland*[26], (iv) *University of Chicago Computational Institute*[27], (v) *National Grid Service (NGS) in the UK*[28] and (vi) *Grid Ireland*[29]. Within New Zealand, the following institutions have been investigated:

---

[24] www.ictregie.nl/publicaties/nl_08-NROI-258_Advies_ICT_infrastructuur_vdef.pdf

[25] www.vpac.org

[26] www.csc.fi/english/csc/overview/organisation

[27] www.ci.anl.gov/

[28] www.ngs.ac.uk/managementboard/

[29] www.grid.ie/

(vii) REANNZ—the company that delivers the KAREN network, (viii) MacDiarmid Institute, (ix) Maurice Wilkins Centre, (x) NZ Synchrotron Group, and (xi) NZ Genomics Limited.

Where possible, we conducted interviews with the leaders of these infrastructures and consulted their founding documents, business agreements and financial models and used our findings to shape NeSI's governance and planned operation. Many governance and management principles were adapted from NZ Genomics Limited since this infrastructure has several operational similarities to NeSI.

An MOU, Funding Agreement and Service Level Agreements were chosen because these have proven to be highly effective at coordinating the efforts and activities of several research institutions in a fair and transparent manner, and all Principal Investors have a familiarity and comfort with this mode of collaboration.

The University of Auckland has agreed to act as the host institution for NeSI and acknowledges that this does not imply ownership or additional rights over the infrastructure.

5.1.2 Ownership and Management of Resources

The equipment purchased by the capital investments made at the Principal Investors will be legally owned by each institution. For administrative and human resource purposes, personnel will be employed by Principal Investors, with lines of operational responsibility vested in the NeSI partnership and with all personnel forming and participating in the NeSI organisation. NeSI will secure supply of its services and workforce via Service Level Agreements with each of the Principal Investors.

The level of commitment demanded from Principal Investors is high, necessitating a commensurately high degree of shared understanding and trust around the NeSI partnership. The common objective is to facilitate deep collaboration—effectively placing institutional needs on par with those of the wider research community—and covers shared planning for services and computing platforms, coordinated outreach, a 'single front door' managed access strategy, and transparency in monitoring performance levels of services provided and of uptake by various communities.

The structure of the NeSI community is shown in the figure 7, and the goals and principles around which the NeSI partnership is based are detailed in the following sub-sections.

5.1.3 Membership and Participation

There are several types of membership within and participation with the NeSI partnership, ranging from direct investment through to elective participation. The four direct, or Principal Investors are the University of Auckland, University of Canterbury, AgResearch/Otago (represented through AgResearch), and NIWA, referred to as the Principal Investors. Two Associate Investor research institutions will contribute investment via an associated Principal Investor. Associate Investors are the University of Otago and Landcare Research. They play a significant role through joint investments and related resource-sharing agreements with AgResearch and the University of Auckland respectively. These types are described below.

Principal Investors

The Principal Investors are those institutions being direct parties to the legally-binding agreement that forms the NeSI partnership. Through joint control of NeSI their common objective is to provide a highly cohesive infrastructure that will serve the needs of all NZ research institutions with focus on strategic research communities.

The Principal Investors, through their governance role on the Board, have opportunity to allow new partners to be added, and existing partners to retire or change. It is not envisaged that investment levels or composition and number of Principal Investors will change. The partnership has a mix of universities and CRIs, and four partners is a reasonable, manageable number for direct governance.

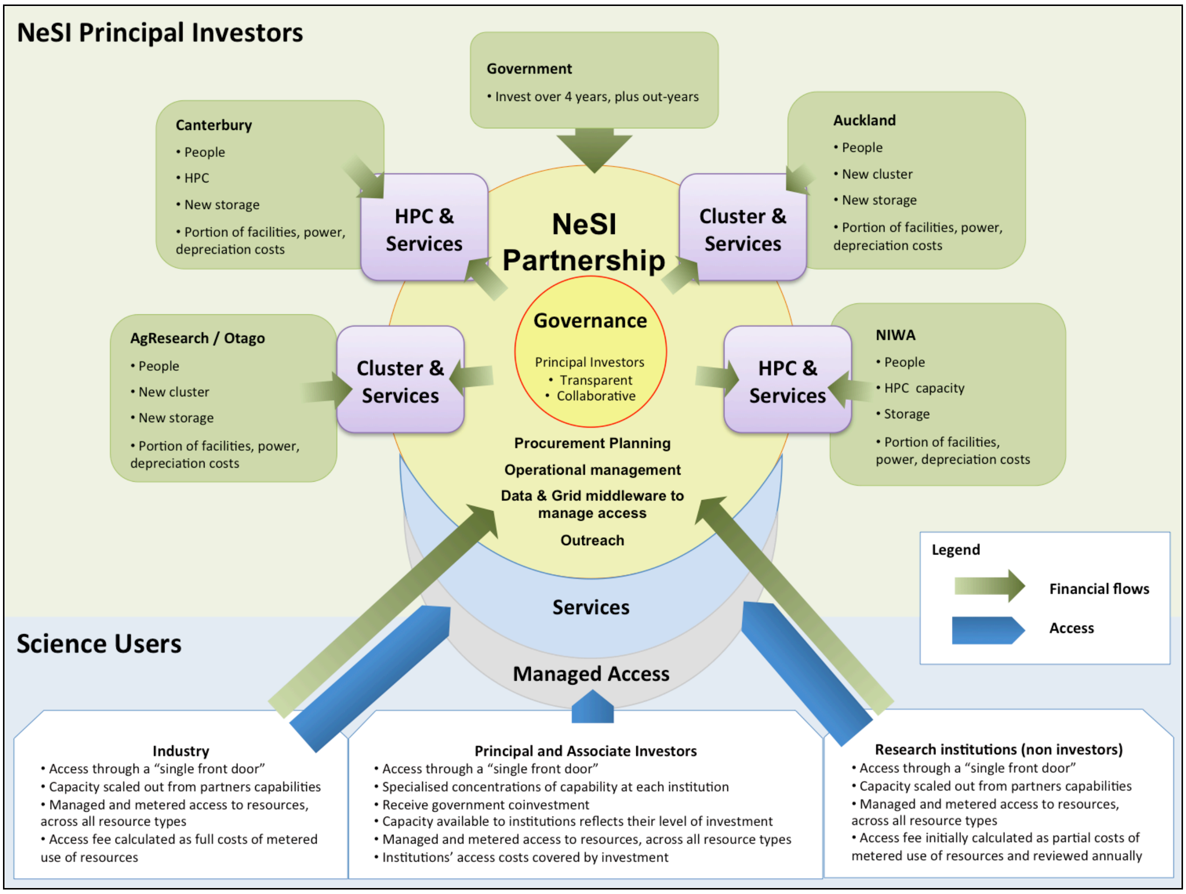

Figure 7. The NeSI community: **Principal Investors and Science Users**: Principal Investors are shown in green and comprise the government, University of Auckland, University of Canterbury, NIWA, and AgResearch/Otago. Facilities are shown in purple. Managed Access to most Services is via a 'single front door' available to (i) Industry (ii) Principal Investors in relation to their investment and (iii) all Research Institutions via an Access Scheme.

Associate Investors

Associate Investors are co-investors into the partnership, with their investment channelled through a Principal Investor via side agreements between these parties. This establishes the partnership as a small tight-knit core of Principal Investors, while allowing for participation and coordination of a wider pool of resources. The expectation is that Associate Investors will place their resources within the control of—and collocate them directly with—their associated Principal Investor, thereby contributing directly to that Principal Investors resources within the partnership. Any rights and obligations accruing to the Principal Investor also propagate to an Associate Investor.

Sector Affiliate

In order to continue with the existing model developed within BeSTGRID, Sector Affiliates describes an option for elective participation. Sector Affiliates are institutions choosing to coordinate their resources within the NeSI infrastructure, with the benefits of providing a unified set of approaches allowing easier consumption of their own resources and their contribution to an overall increased scale of resource to the wider sector. In practice this would mean adopting the operating principles of the Charter (Appendix 5), and implementing common middleware and service configurations. To recognise the overhead to the NeSI Principal Investors of supporting this approach, these Sector

Affiliates will be required to contribute a commensurate level of FTE to the initial build and ongoing support and to commit spare compute capacity back to NeSI for education access.

5.2 Governance and Management

The following organisation structure will be constituted for Governance and Management responsibilities:

1. A **Board** responsible for strategy and policy and approving major initiatives and investments, including the allocation of resources centrally and within each institutions, and other executive decisions
2. A **Director and Management Team** for planning and overseeing day-to-day operations of facilities and development of services
3. A **User Advisory Group** to engage with the Management Team on day to day operational issues and provide input into service development and management

Figure 8 below shows the relationships between each of these three bodies.

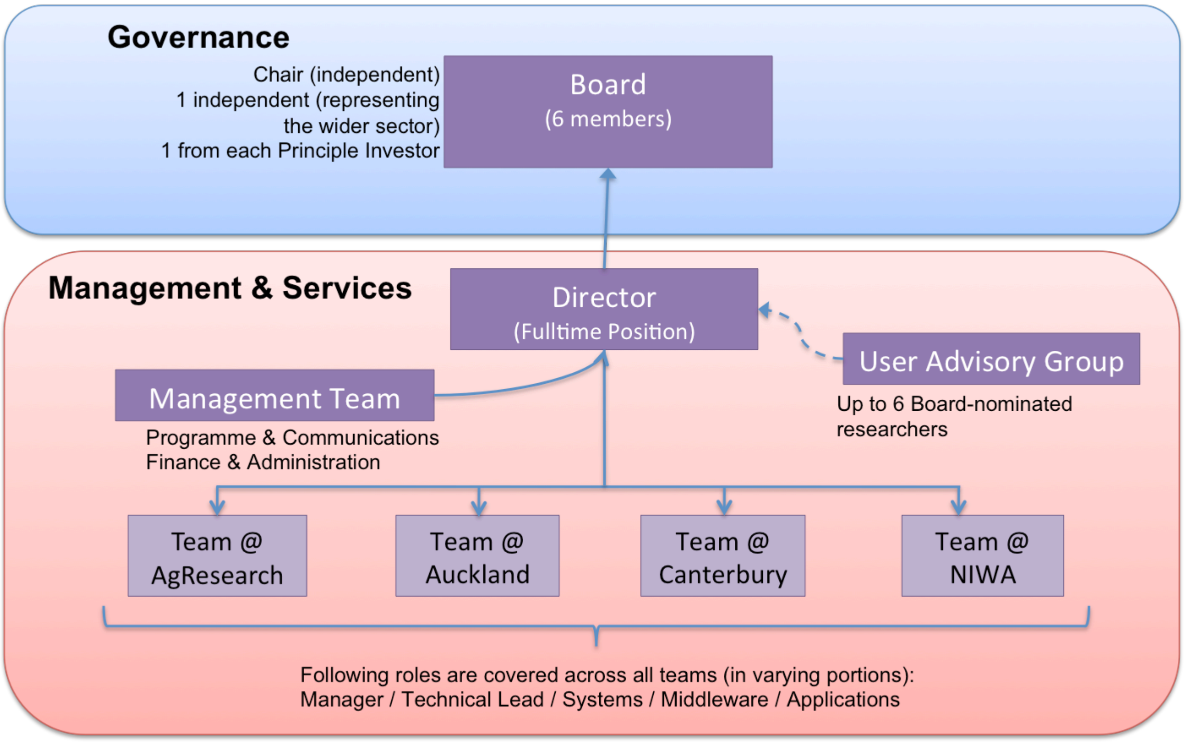

Figure 8: **Organisational Structure** of the NeSI partnership showing Governance (blue) and Management & Services (red).

5.2.1 Board

Governance and executive control will be by means of a six member Board consisting of individuals with broad experience in research, HPC, eScience or governance. Each of the Principal Investors will occupy a single seat on the Board. The Board will also include an independent Chair (chosen for skills and experience in chairing) and another independent member with expertise in the field representing the non-investors (i.e. the research sector at large). Associate Investors will be represented through their Principal-Investor partners. The independent members will be selected after consultation between the partnership signatories. Appointment of the Chair must be approved by all four of NeSI's Principal Investors.

Were the Board to make a decision that the Host believes is at odds with the Funding Agreement, the Host will have the opportunity to appeal for resolution, in contractual terms. This appeal would only be in light of the Host delivering on the Funding Agreement, and will not give the Host veto rights over the normal decision making of the Board that is in line with this Investment Case and the Funding Agreement.

Board Responsibilities

The key responsibilities of the Board include:

- Ensuring the infrastructure remains focused on the needs of the NZ research community, within the Scope of Activity defined by this investment case and associated agreements, and within the constraints of the budget
- Accountability to the Principal Investors, funding bodies, and the wider sector for the partnership's performance
- Advising on the selection and appointment of, and reviewing the performance of the Director and Management Team
- Review of longer-term trends in the use and uptake of the infrastructure and approving investment plans
- Development and review of policy for access
- Responsibility for finances and funding, and specifically for ensuring a continuing funding-base
- Coordination with government agencies on policy and funding implications

The Board should meet at least four times each year, in person where possible. They may meet more frequently in the first year, as needs determine.

The Board may establish committees to provide advice, develop strategies, policies, and plans, and review performance. Board decisions will be by consensus with recourse to a dispute resolution process. The Board retains authority over partnerships involving NeSI with third parties, which includes oversight and authorisation of any allocation of rights, co-investments with third parties, and other such relationships that may dilute its shared objectives, and the rights and obligations of its investors, change it's risk profile, or in any other direct way impact on NeSI.

5.2.2 Director and Management Team

Day-to-day operation of the infrastructure will be the responsibility a Management Team, whose members include a full-time Director (who will report to the Governance Board), an administrative team at their local institution, and a staff member located at each Principal Investor responsible for service management in line with their Service Level Agreement. The administrations team consists of two support staff who will be working with the Director on operational matters, being a programme manager with a joint focus on outreach and communications, and a business support staff member dealing with finances and administration. See Figure 9 below.

The role of the Director is to design, develop and implement strategic plans in a cost-effective and timely manner. The Director is also responsible for directing the day-to-day business of the organisation, including leading the Management Team, developing business plans in collaboration with the Board for the future of the organization, ensuring Service Level Agreements are adhered to, and prioritising various development projects. Specialist teams will be created and tasked by the Director to carry out work on specific areas of the infrastructure.

Appointment of the Director must be approved by all four of NeSI's Principal Investors. The Director is accountable to the Chair of the Board and reports to the Board on a regular basis at each Board meeting. The Board will review the operations of the infrastructure to ensure that community needs are assessed, appropriate provisioning plans are made and fair access schemes are provided.

The Management Team will meet regularly, to support the Director in development of plans, review all Service Level Agreements, and track progress on initiatives including service developments and operations. Note that the focus, operation and membership of this Group is similar to the current BeSTGRID community, which has an effective long term governance and management structure running since 2006 across 9 institutions, and whose governance and management practices will be reviewed to inform the establishment of this Management Team.

Principal Investors will negotiate with the Director and provide recommendations to the Board on committing suitable staff to the Management Team and each Principal Investor's team. Candidates should meet the partnership's requirements for capability to fit the roles agreed in the SLAs and be full-time, or as close as possible. The resulting distributed team will be coordinated centrally on shared goals, and be required to align institutional approaches to agreed standards and platforms.

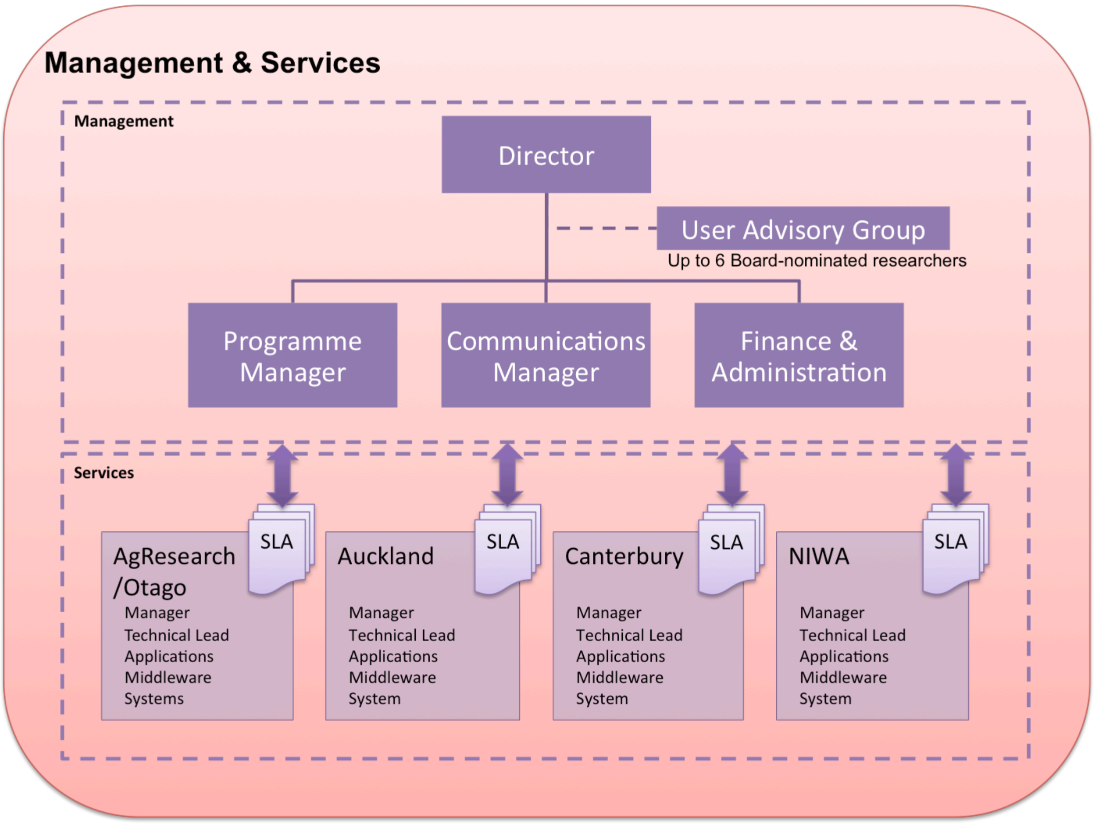

Figure 9: **Management and Services**: The Management Team includes each Principle Investors' team Manager, who is responsible for managing their Service Level Agreement. All members of each team are members of the Services Team.

5.3 Business Processes and Services

The core business of the NeSI partnership is to coordinate capital investment in HPC and eScience efficiently and operate it effectively for access by the NZ research sector. The key business processes and the related services are depicted in Figure 10, and further described below.

Access Policy: Policy outlining how the access scheme should be constructed and operated, incorporating definition of the Investor, Commercial, Merit[30], and Educational Queues, along with

---

[30] Refers in the general case to NZ publicly funded research, with notable exceptions.

definitions of who qualifies for access to each queue[31], and the appropriate prioritisation between these queues. Responsibility for developing and applying this policy rests with the Board.

Strategy & Investment Plan: The business strategy and annual investment plan. This is developed by the Director, and approved by the Board. This feeds into Procurement Coordination and Resourcing.

Resourcing: The resourcing function focuses on mapping the Access Policy to appropriate resources, based on analysis of resource capacities and historical utilisation, and guides the implementation of the Access Scheme.

Audit & Accounting: This is a crucial function in order to achieve transparency of resource utilisation, especially to ensure the principles of being equitable and inclusive are adhered to, and that the wider NZ research sector feels like the organisation is accountable, and hence participation is desirable. This function will bring together a standard set of measures from across all the institutions providing resources within the infrastructure, and review against the Access Policy and the Strategy & Investment Plan, to ensure operational alignment. Reports on progress related to these aspects of operations will be provided at each Board meeting to inform strategy and investment plans and access policies, and agreed outputs will be available to the wider sector within the Reporting function.

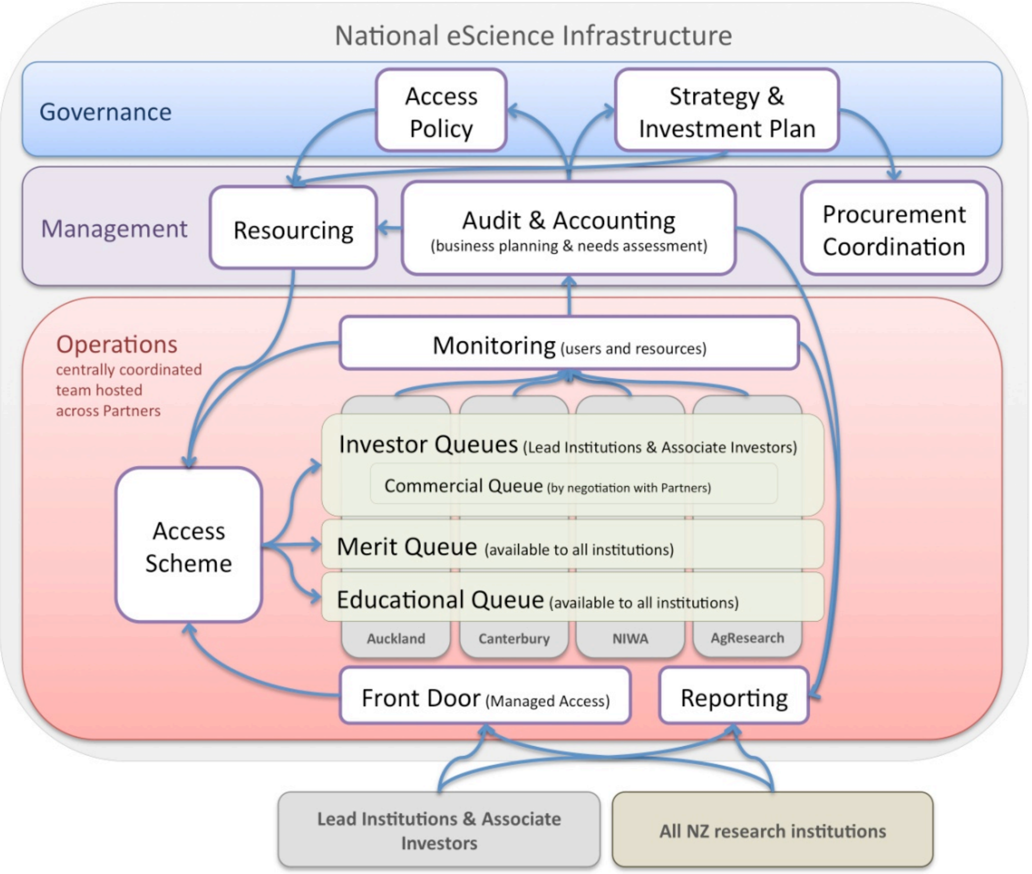

Figure 10. Interactions between Governance and Operation of the Infrastructure. Colours as per previous figure, with the infrastructure in green and partner institutions are shown in grey.

---

[31] The Board will need to establish what happens in the case of developed core funding for the CRIs to inform the definition of Merit.

Procurement Coordination: The Principal Investors and Associate Investors agree in principle to a high level of transparency on procurement processes and decisions, including ensuring that the Board is party to any agreements with vendors executed related to investments allocated by the Board; subject to confidentiality as necessary and at the Board's discretion. This coordination will not extend to control over procurement processes; rather it aims through transparency to ensure that every opportunity to achieve maximum buying power and efficiency is taken, and that the trust between institutions is maintained.

Software: In general, the NeSI partnership cannot take financial responsibility for purchasing of software licenses used across the science sector for specific research needs, but in coordination with the Tertiary Technologies Purchasing Committee it may help facilitate the joint serving and managing of these licenses, so that they may be pooled for use across the infrastructure. Where software is of wide applicability, the partnership may choose to purchase licenses.

Access Scheme: The Access Scheme is the tangible artefact implemented across the sites providing services. It implements the principles and priorities of the Access Policy, and maps these onto institutional resources, inside the following Queues:

> **Investor Queues, Commercial Queue**: The Investor Queues are available to the Principal Investors and Associate Investors, as a reflection of the direct investment each has made in the infrastructure. Within these capacities each institution is able to allocate and prioritise access, for their own use of a research or commercial nature, without prejudice or the need for consultation with the Board. This includes allowing the free exchange, with capacity at another institution. Partner institutions can choose independently how to utilise their contributed capacity. Unused compute cycles revert to the Merit Pool, which can then be used to run merit jobs with a lower priority.
>
> **Commercial (non-research) uptake**: Commercial uptake of the computing platforms will be encouraged. Most commercial HPC users are likely to utilise the non-NeSI portions of capacity (i.e. Partner capacity) at one or more of the investing institutions, with revenue flowing to these institution providing the capacity. If a commercial application makes use of NeSI-funded capacity and/or services then these must be fully costed and reimbursed back to NeSI. There is already commercial use of the Canterbury, AgResearch and NIWA facilities, so by adding in additional capacity we create the opportunity for this uptake to increase. In either case, the NeSI team will work with NZICT and NZTE to promote the use of the infrastructure to the commercial sector, and with research funding agencies and research institutions to encourage uptake among the CRI and University communities.
>
> **Merit Queue**: The Merit Queue reflects the capacities available for access by all institutions, based on the definition of Merit within the Access Policy. This queue makes no distinction as to the affiliation of the proposers, who may come from the commercial or research sector, but it does make distinctions based on the assessed quality of the research proposed, and its fit with sector priorities for research and economic growth. Standard production applications from either the commercial or research sectors will not be eligible; the research proposed must have demonstrable science merit. To ensure that benefits accrue to New Zealand, priority will also be given to NZ-funded research. Other non-private funding such as offshore publically-funded research or research funded by charitable trusts may also be given some priority.
>
> **Education Queue**: Spare compute cycles (unused from the above queues) will be made available for educational purposes. It is envisioned that such access will be allocated to postgraduate and early career researchers, who otherwise might have difficulty qualifying for Merit Access, for example. This queue will have a lower priority than the Merit Queue, and hence queue wait times would be longer, though it will provide an opportunity for access to advanced computational infrastructure beyond any institutions means to provide.

Further refinements to the Access Scheme will be required to manage and balance computational workloads. For example, long-running jobs may be given a lower priority to ensure that no single researcher monopolises or captures any system.

Monitoring: All resources within the infrastructure will be monitored on a similar and shared basis. Management will design a set of metrics and measures to ensure audit and accounting functions can be routinely supported.

Reporting: Reporting will be carefully constructed to allow transparency of operations and allocations, to ensure that the wider sector can review the application of the Access Policy, and be informed on the performance of the infrastructure and its services.

Front Door / Managed Access: This key function supports the unified design for accessing infrastructure services. The principle of a single front door reinforces the strong desire to have a single integrated set of services presented to the wider research sector, and provides the opportunity for monitoring, auditing, and accounting to operate in a shared and similar manner.

The above business processes and services represent those of primary importance in operationalising the good governance needed to facilitate a nationally focused infrastructure service. They are not suggested to be comprehensive, though provide a strong indication of the basis on which the Board will establish the organisation and infrastructure services, and are the formative basis for an initial business plan.

There will be additional issues that need to be resolved by NeSI governance as the project becomes operational. A brief list of issues appears in Appendix 6.

5.4 Intellectual Property Rights

As a general guiding principle, any service that has been developed and can be re-used for no cost will be made available freely, across the research and commercial sector. This will act in the public good, building up NZ national capability in science infrastructure.

The following five clauses will apply across the NeSI partnership:

1. Any and all Existing Intellectual Property Rights and any modifications of those Existing Intellectual Property Rights shall remain the property of the original owner.
2. Subject to clause 4 (below), any Intellectual Property Rights (“IPR”) created by any Collaborator or NeSI-Affiliated Persons in the course of providing a NeSI-Affiliated Service shall be owned by the Collaborator and be subject to the policies of the Collaborator, unless some prior agreement over-rides this arrangements; for example some community-based computer programs require improvements in the code to be shared back with the broader community.
3. If NeSI-Affiliated Persons from multiple NeSI Collaborators are involved in any specific piece of work that generates IPR in the course of providing a NeSI-Affiliated Service, such IP will be subject to a separate agreement specifying the ownership share. In the event of a dispute, the Board shall determine the ownership of such IPR.
4. IPR generated as a result of services provided by a NeSI Collaborator for a customer in the course of providing a NeSI-Affiliated Service shall vest in the customer unless agreed otherwise in a specific research or collaboration agreement.
5. IPR generated in the course of providing a NeSI-Affiliated Service and owned by a Collaborator will be freely available to NeSI for the purposes of NeSI work subject to such use not being detrimental to any private commercial interest of the Collaborator recognised in the relevant Service Agreement.

5.5 Assessment of Effectiveness of the Infrastructure

Several different strategies will be used to determine effectiveness of the infrastructure. These can be broadly grouped into two sets: those that measure the uptake and efficiency of the compute platforms, and those that measure the effectiveness of the infrastructure at the level of national impacts.

5.5.1 HPC performance metrics

We assess the efficiency and utilisation of the HPC platforms using the information statistics gathered from operations management. Three distinct questions are addressed: (a) how is it being used, (b) how efficiently is it being used, and (c) how good is the service level?

Metrics to answer **how the infrastructure is being used** (a) and to **assess efficiency** (b) include:

- Measures of resource requirements, including core count, queue time, processing time, memory utilisation and input / output rates, data input / output rates, disk allocations / occupation / access rates.
- Measures of user activity segmented by their institutional and research community affiliations, scales of use, and applications.
- Tracking of improvements to scalability (e.g. as measured by speedup) of codes (over time).

Profiles will be developed and publicly reported summarising all the above measures for each institution and research community.

Metrics to describe the **level of service** (c) include:

- Overall machine availability (by time, by platform).
- Batch queue lengths and times.
- Slowdown i.e., (runtime + wait) / runtime: reported by runtime, by job size (cores) and execution time (or queue).
- Number and level of severity of system incidents, mean time between failure.[32]
- Systems support issues: time to respond to queries, time to resolve, outstanding queries.
- Scientific programming support: time to implement (including parallelization) and optimize new and/or existing user codes.
- Development, growth and sustenance of user-communities: number of new users and applications, percentage of off-site and outside-of-community users who seek engagement, retention of user communities.

5.5.2 Assessment of national benefit

These assessment criteria are linked to merit-funded science and will use the output information so generated, as well as capturing details of the science enabled from the researchers, and feedback from them on their experience.

- Successful grant application that include funds for HPC or eScience services including nationally and internationally.
- Satisfaction and impact survey of the user community (annually undertaken).
- Development of all IP that has benefitted from this investment (across all platforms and services).
- Number and Value of commercial contracts obtained that use services.
- Number and Value of 'collaborative' projects within and between CRIs and Universities utilising services.

---

[32] www.hector.ac.uk/about-us/reports/annual/2009.pdf

- Project description and brief annual project report required from all users.
- Journal Publications that used the infrastructure to obtain results.
- Yearly reports from the Board, Management, and Operations team outlining impacts, breakthroughs and publications.
- Media statements, both from within the NeSI partnership and from associated user communities.
- A series of more detailed case studies to show effectiveness of engagement and support of specific research projects.

5.5.3 Government evaluation of the investment

Government, via its agent the Ministry of Science and Innovation, will perform an implementation review following the first year of operation of the infrastructure, at a date agreed in the Funding Agreement. The Board will provide it with the information required in a timely manner, according to terms laid out in the Funding Agreement. A full review will be conducted at year 4, prior to continued investment.

Based on the accumulated evidence and external evaluation, a more detailed review will be conducted in the fourth year to determine:

- The effectiveness of the NeSI partnership.
- The impact of the investment.
- The change in behaviour of the institutions and funding bodies—towards a more coordinated and collaborative model.
- Up-to-date opportunities for renewal of the investment.

At this time, the infrastructure will be aligned again if needed with emerging and evolving needs, and with current trends and new technologies available.

5.6 Risks and Mitigation

The following set of risks and mitigation strategies have been identified and are listed below in Table 4, decomposed as organisational, technical and operational risks. As can be seen, most risks are low because this project builds upon the significant expertise of the partner institutions, all of which have been successfully running large-scale computational research infrastructure for a number of years. The only risks with greater than Low impact are those relating to loss of key staff and tension among researchers over the allocation of resources.

Table 4. **Risks**, a summary of risk, likelihood, impact and mitigation strategies to be pursued.

Organisational risk		
Risk	**Likelihood/Impact**	**Mitigation**
Failure of the NeSI partnership structure to work effectively	Low/High	Establishing appropriate agreements, including the NeSI MoU and Service Level Agreements. Independent & unified (national) branding of infrastructure services. Appropriate overview by government (funding) agencies.
Acceptance and adoption by the science community	Low/High	Develop transparency in governance and management and seek active participation. Align the capabilities, capacities and support with needs of the research community.
Change to (reduction in)	Low/High	Investors underwrite the risk for their individual

government or partner contributions		components of the system, and commit to provision through Service Level Agreements.
Changes in government policy and priorities	Low/high	Work closely with key agencies to understand and respond to priorities.
Failure to achieve sufficient user revenue	Low/Medium	All at risk capital is isolated in the budget so any shortfall will not cause the partnership to fail. Work with users to explicitly budget for services, and work with key funding agencies to recognise costs.
Poor computational performance for users	Low/High	Extensive benchmarking of computational performance of relevant user software in the tendering and acceptance testing phases of acquisition. Team established capable of supporting migration and refinement of new and existing codes onto suitable platforms.
Projected usage not met leading to unused capacity	Low/Low	Current user demand and national survey all point to a very large, and mostly unsatisfied, demand. Future investments are made once the overall usage, demand and available funding have been assessed. Unused capacity is provided to Educational community, to train new generations of researchers.
Ongoing investment is not met by partners	Low/High	Secure investments to ensure capital is available. Strong engagement with senior management teams at all partner institutions.
Technical risk		
Performance risk - systems/hardware not fit for purpose	Low/High	Ensure strong and ongoing alignment with the needs of the research sector. No experimental or first of type to be included in infrastructure until externally validated. Board responsible for guiding procurement of fit for purpose systems. Reliability to be assessed prior to purchase.
Data Transmission risks	Low/Medium	Work closely with REANNZ to ensure close integration with the KAREN research network and meeting the traffic demands of the research sector.
Implementation delays	Low/Medium	All partners have strong experience in planning, installing and operating large computational systems and services. This combined experience will be leveraged at all stages.
Computational infrastructure (e.g. power) failure	Low/Medium	Partners contract to provide the appropriate support with recovery plans via Service Level Agreements.
Operational risk		
User requirements not met, computationally or for robust services	Low/Medium	Extensive evaluation of user needs prior to purchase and in configuring resources – e.g. matching/balanced data storage. Many applications will be able to operate on a variety of platforms although in some cases with lower

		efficiency. Data systems to include redundancy. Fail over capability to be provided.
Users not satisfied with resource allocation	Medium/Medium	Create transparent processes for resource allocation, taking account of government priorities.
Inadequate system support	Low/Medium	Partners undertake to provide adequate support and end-user engagement via Service Level Agreements.
Loss of key staff	Medium/High	Build in redundancy by sharing knowledge, skills and practices widely. Investors undertake to provide appropriate support and training, and recruit at the appropriate level.
Users make limited use of the facilities	Low/Medium	Investors provide support and training for systems and applications to up-skill users.
Security	Low/Medium	Best practices used throughout to mitigate risks of compromised systems, and appropriate levels of data security and identity management used.
Natural hazards	Low/High	The distributed nature of the system provides some diversity. Multiple communication links to increase resilience for nationwide users.

6 FINANCIAL DETAIL OF THE INVESTMENT CASE

The NeSI partnership will coordinate capital investments between four Principal Investor institutions with capital grants made for establishment of essential capacities, and each Principal Investor will deliver services under a Service Level Agreement established within a contract held with the host institution. Two Associate Investor research institutions will contribute investment via an associated Principal Investor. The institutions are:

Principal Investors:

- The University of Auckland
- University of Canterbury
- AgResearch/Otago, led by AgResearch Limited
- National Institute for Water and Atmosphere Limited (NIWA)

Associate Investors:

- University of Otago, through AgResearch Limited
- Landcare Research Limited, through The University of Auckland

The total costs of the infrastructure are estimated to be $51.8 million over the first 4 years (includes investment and income), with government contributing $27.5 million of total funding required [redacted] Operating costs[33] are part of both co-investor and government investment; as is commitment to reinvest capital. Figure 11 breaks down the investment by year and by source.

[redacted]

Figure 11. **Breakdown of the total cost** of $51.8 million, showing contributions by government, the investing institutions, and the wider sector (merit access and commercial fees).

The investing institutions contribute [redacted] over the 4-year period [redacted] capital and [redacted] operating costs)—this is shown in red in Figure 11. Over the life of the contract the investors contribute [redacted] of the investment, and the government [redacted]. A fixed portion of all resources will be assigned as priority for the investors (in proportion to investment). The remaining

---

[33] Operating costs include establishment and management; personnel and administration, and operating depreciation related to all capital assets.

resources will be dedicated to the wider science sector. Figure 12 depicts the range of annual income over the initial period of four years, divided into funds from government, institutional investment and access fees. At risk funding will not be disbursed until it is received.

[redacted]

Figure 12. **NeSI annual income ranges**, split into revenues from government, partner institutions and at risk income from access fees.

The NeSI partnership will be structured over an initial contractual period of 4 years, with review and renegotiation of government and institutional investments scheduled in year 4, reviewing the success of the infrastructure to that point. This infrastructure is expected to be sustainable only with ongoing government and institutional co-investment, to overcome the high costs of the HPC platforms, and to enable the sector-wide coordination required to build, deliver, and support HPC and eScience services. This infrastructure is recognised in the 2010 Large Scale Research Infrastructure Strategic Plan, produced by MoRST in consultation with the research sector, as being a high priority for long-term investment. The expectation is that at each renewal, and for a subsequent 4 year term, each investing institution will renew their investment into the infrastructure, and this reinvestment will be complimented with appropriately justified investment from government.

The proposal is sensitive to the following financial factors:

- Demand forecasts for non-commercial (publicly funded) services
- Demand forecasts for commercial services

As shown at the top of Figure 12, at-risk revenue expected to be generated by the NeSI partnership is estimated at [redacted] million in revenue from the merit-based allocation scheme and [redacted] million from charges on commercial services provided in support of commercial research contracts over the 4-year period. The merit-based allocation scheme covers [redacted] of costs for accessing a merit pool of resources, with the remaining [redacted] of costs coming from science users, generating the [redacted] revenue referred to above.

This merit access pool is available across the resources that are dedicated to the wider research sector, including research led by commercial organisations. The merit access pool recognises full cost funding mechanisms will need time to adjust as the new scale of resource capacities available through the NeSI partnership will place increasing demands on research funds. It is expected that the renegotiated proposal at year 4 will start to recognise a decreasing proportion of this funding being sourced via the NeSI partnership as pricing signals are recognised through the funding mechanisms—from both government and institutions—and are adjusted to support the full cost of accessing such infrastructures.

6.1 Financial Assumptions

They key assumptions within the financial model are:

1. The initial period of the NeSI partnership is 4 years, with a governance and funding review and reinvestment proposal developed by year 4
2. The Government appropriation for Research Infrastructure is a base line appropriation, providing a stable funding environment over the coming 10 years.
    a. MoRST has developed a Large Scale Research Infrastructure Strategic Plan indicating high priority new investment areas over the next 5 years. Currently MoRST has signalled that within this strategic plan there is planned ongoing commitment to national HPC and eScience infrastructure, as the highest priority research infrastructure investment opportunity (alongside KAREN, on which it depends).
3. The following costs are shared across all investors into the NeSI partnership, including government and the investing institutions, in proportion to their investments in capital equipment:
    a. Capital costs including operational depreciation on capital equipment
    b. Personnel contracted through Principal Investors
    c. General operating costs
4. Principal Investor commitments lock in institutional capital and operational budgets for 4 years, incurring an opportunity cost to the institutions in terms of other investment priorities over this time frame.
    a. Capital investments are into fast depreciating assets that are a short-term investment with little to no residual value after 5 years
    b. The budget model describes commitments by institutions over the first four years in terms of continued funding of computational facilities. Principal Investors will review commitments at year 4, and reinvest dependent on the success of the venture.
    c. Depreciation accrues at the host institutions. Any unused, government-funded depreciation will be used to offset further investment if this is undertaken in year 5.
    d. The expectation on behalf of government is that depreciation costs disbursed to the collaborators will be reinvested into the infrastructure—in essence then government see this as funding of a ‘capital reserve’ that can be applied to help sustain the infrastructure into the longer term. Provision is made in the budget for asset replacement at the conclusion of the first 4-year period (phase 1). The strategic intention of the investing institutions is to make provision for reinvestment into phase 2 of the infrastructure (years 5-10—shown in Appendix 7).
5. The NeSI partnership will solely meet the costs of:
    a. Governance and secretariat
    b. Management Team employees (2.6 FTE), within the following roles:
        i. Director
        ii. Programme Management and Communications
        iii. Administration and Financial Management
6. Commercial pricing schedules are set at least 100% of the full cost of providing services.

7. The at-risk portion of the NeSI budget, arising from the take-up of merit access, will be reinvested in people. This will keep service levels constant as the user-base expands. At risk income will not be spent until such time as it is in hand, so that the infrastructure does not incur an operating deficit.
8. A 3.5% cost of living adjustment is incorporated into salary-related costs on an annual basis.
9. The merit access scheme aims for a balance between stimulating merit access and sending a price signal to funding providers and researchers in order to provide increased capacity and incentive while ensuring control of demand.
    a. In approving access for any researcher via the merit access scheme a key principle is that the government should not pay twice, i.e. as the government invests in the partnership to provide HPC capacities to the wider sector it should not be asked to fund researchers for the full cost of accessing said HPC through other funding sources.
    b. To allow a smooth transition in funding arrangements, there will be no displacement of existing funding arrangements. So if research funding is currently or in the future awarded with full cost of HPC and/or eScience included, then all concerned will be required to honour this commitment.

6.2 Financial Risks and Mitigations

Some basic financial risks are common to many projects, namely foreign exchange impacts on the cost of equipment and associated consumables (all are imported), and power or other major operational costs. Such risks are inherent to undertaking research and are no different for the partnership than for individual research institutions embarking on investments in HPC.

In addition, there are specific risks in establishing an infrastructure of this nature; namely the difficulty in initially activating collaborators to form such a joint infrastructure, and the operational challenge of then exploiting and sustaining the investment. Each of these is addressed in turn.

Activating coordinated investment

Direct financial contributions to a shared research infrastructure are difficult to incentivise, because institutions find it hard to monetise their investments when the benefits are distributed beyond the investing institution. There is a cost associated with coordination and providing a national service. In addition, it takes time to reap the benefits of this kind of investment.

Because many institutions are not prepared to add another membership fee (the likes of KAREN) for access to research services, the investors in the partnership have decided to activate coordinated investment themselves to begin the process of developing a national infrastructure.

This initial coordinated investment is only possible with at least equal government co-investment to provide the incentives for coordinated investment, and to cover the costs of coordination and national service delivery.

In the specific instance of HPC and eScience, the risk of beginning this coordinated investment is further complicated by the exponential growth in computing requirements in research disciplines with computationally and data intensive computing needs, creating ongoing pressure for greater investment.

The investors in the partnership, many of whom have been investing for their own institutional requirements for some time, recognise the need to coalesce their resources into joint and larger-scale capacities, and to scale these to meet demand across the sector over the longer term.

The considered response to this concern from the investors and the sector has been to:

- Establish a small and tightly aligned leadership group of Principal and Associate Investors, that is willing and able to coordinate significant resources and strategic directions to form the partnership under joint control and with a workable budget from the outset. This forms the nucleus of the capability, which is scaled out to meet further sector wide needs, thereby avoiding the situation where the sector could be asked to fund a budget shortfall.
- To address these sector-wide needs and support researchers' access to the infrastructure, allocate capacities for merit-based access and require partial funding contribution (20% of costs) rather than annual subscription or membership fees per institution. This allocates costs directly to researchers rather than their institutions and the sector bearing these costs indirectly and in some cases disproportionately. The establishment of the merit access scheme will be in consultation with appropriate funding bodies, and hence these costs will be made clear to them, allowing funding schemes to be adjusted over time to recognise and utilise the new infrastructure, reinforcing the existing full-cost research funding scheme.
- Recognise the opportunity cost to investors of committing and coordinating significant strategic capital investments and exposing related strategic capabilities through this joint venture. Funding is provided to investor institutions slightly beyond their initial investments to cover the costs of scaling up their coordinated investments in resources for the wider sector at required service levels and offering greater levels of service than those required to run similar infrastructure within a single institution.

The initial success of these approaches is evidenced by the membership of the joint venture representing almost all of the major institutions advancing HPC and eScience in New Zealand, and by the wider sector support being received for this proposal from non-investing institutions who can see the benefits the infrastructure will provide to their researchers—see Sections 2.6 and 3.5.

Exploiting Capital Investments

The ability to operate and leverage the capital investments in research has historically presented challenges, partly due to the difficulties in sustaining operational funding in support of step change increases to larger capacities. This pressure on operational funding often cannot be accommodated within existing full-cost funding mechanisms that have been scaled for previously established capacities. Suitable periods of readjustment are required before a new steady state is established.

The Transition Working Group has carefully sized the capital investment and balanced its investment with appropriate personnel that will support the sector in exploiting the new capacities. It is expected that the demand for support and consultancy services will grow over time, as the growth in access to and use of resources puts greater demand on the resources, and requires more considered optimisation and tuning to get the best return on the partnership's investments. The allocation of staff, roles, and responsibilities will be carefully monitored and adjusted by the Board to address any issues in fully exploiting the capital investments.

6.3 Liabilities

The liabilities due to significant capital equipment acquisitions relate to their ownership, depreciation and other operational costs, and requirements to operationalise investments to deliver services to users.

The equipment procured through this funding will be owned by the investing institutions. These institutions have agreed to an initial sizing and distribution of equipment in a manner that also aligns with their own capabilities and expertise, although at increased capacities closer to those required to meet indicated demand across the research sector. This alignment between research infrastructure priorities of NeSI and investing institutions increases the potential utility of assets, reducing the investment risks.

Each institution will retain liability for the portion of the investment that they are contracted to provide. Joint liability will only be held for clearly identifiable joint activities and decisions. In both cases, liabilities will be limited to the extent of the institution's investments into the infrastructure.

The potential (operational) risk is reduced further for major categories of equipment through the partnership allocating funding to purchase agreed resource capacities from each investor over the term of the MoU Funding Agreement and SLAs, as they fit with NeSI's requirements and meet ongoing contractual obligations defined within each Service Level Agreement.

The liabilities accumulating to these institutions from operating HPC equipment come in the form of depreciation and additional personnel to maintain, support, and exploit the equipment. The institutions will align their own operational budgets in proportion to the capital investments they will each make, and apportioning costs from within the government budget will cater for the scaling up of operating budgets required by the increased scale of services and equipment.

6.4 The Proposed Capital Investment

The proposed capital investment is [redacted] which is spread across each year of the first 4 years of the investment to recognise that capital investments will be in the form of finance leases in almost all cases, to provide lower capital risk and preserve some upgrade options. The collaborating institutions are committing [redacted] million over the first 4 years to support the purchase and lease payments on the computing asset necessary to support NeSI with government funding meeting the shortfall of [redacted] million.

This new equipment will be combined with existing eScience equipment owned and operated by the Principal Investors, and in some cases by other organisations with a preference to coordinate with NeSI. The capital equipment will be operated by NeSI and supporting personnel employed by the collaborating institutions.

Investing institutions have agreed within NeSI's Transition Working Group on the allocation of capital investments. These agreements support investment in the types of equipment required to successfully meet the wider research sector needs, which were established through a *National Survey of HPC and eScience* needs, as discussed in Section 2.6, and are described in Section 2.7. Working closely with the partner institutions, NeSI's Board will coordinate and approve future investment plans, reflecting on the service delivery success and take-up of investments at institutions and reprioritising investments on an annual basis.

The major exception to this principle is in funding the largest capital equipment purchases (i.e. those at Canterbury (NIWA's HPC facility is already purchased). These kinds of HPC platforms are best purchased in single, large investments due to the premium for such specialized architectures and price points used by vendors. Others, such as the Super Clusters and Storage, are possible to build up over time. For large capital acquisitions, when smooth cash flow is necessary to meet the even funding from government (as is the case here) a vendor finance lease is the most cost-effective option. Such leases typically include a flexible upgrade path to accommodate advances in computational power that will occur during the life of the asset, thus helping to future-proof the investment. Finance charges within such a lease are also likely to be smaller than the equivalent interest payments incurred for a cash loan.

During each 4-year term a portion of capital investment may be retained and reinvested at the start of the subsequent term, or earlier, as deemed appropriate by the Board. This provides flexibility in allocation of capital to meet any variance in demands placed on specific types of equipment.

The proposed capital expenditure for the term of the initial MoU, Funding Agreement and SLAs is shown below in Table 5 below. A full 10-year forecast of capital expenditure is given in Appendix 7.

Table 5. **Capital Expenditure** on new equipment for NeSI (initial 4 years).

[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]					
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]					
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]					
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]

Table 6 below adds in further details of the planned Capital Equipment (core-counts and storage) that will be acquired on behalf of NeSI each of the four major computational platforms. Each figure represents the new investment in that year, not the cumulative capacities. Note that NIWA's capacity figures in the table are accurate rather than an estimate. They represent what NIWA have already procured, built up, and now operate inside their HPC facility.

These figures represent conservative estimates of the scope of Capital Equipment that could be achieved with the investments, not a specification to be used to refine the annual budget. The priority for the Investors will be to maximize the computational resource that can be purchased for each investment cycle, recognizing that vendors' responses to any Request For Tender will vary considerably for different hardware specifications. Competitive tendering will be used when planning purchases to ensure that value for money is obtained from the vendors.

Section 4.1 above details the current state of the International HPC market and the rationale for not relying on it for the country's HPC needs.

NeSI aims to build balanced systems with a range of performance and behavioural characteristics—but that performance balance is in no way captured by the simplistic capital resource table above. Confidence can be garnered by looking at each of the Investors track record and experience in delivering systems that achieve the needed performance balance at the time we each made our respective investments.

Investments of this nature are very infrequent, and overlay complex and rapid technology evolution. This makes any specific previous investment a very imprecise indicator of future cost breakdowns and specifications. To offset the risk, some equipment will be acquired up front, to provide an immediate large increase in capability, but some will be purchased incrementally, to take advantage of emerging improvements in technology. Some of the systems, such as those at Canterbury and NIWA, are only viable to purchase in single large investments due to their premium for such specialized architectures. Others, such as the Super Clusters and Storage, are possible to build up over time.

Table 6. **Capital Equipment Indications**, showing the (likely minimal) level of capability at each of the four sites, and for each HPC platform type, per year of planned investment. Storage capacity is also shown.

	[redacted]		[redacted]		[redacted]		[redacted]	
	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]								
[redacted]	[redacted]							
[redacted]								
[redacted]	[redacted]							
[redacted]	[redacted]							
[redacted]								
[redacted]	[redacted]		[redacted]		[redacted]		[redacted]	
[redacted]	[redacted]		[redacted]		[redacted]		[redacted]	
[redacted]								
[redacted]		[redacted]		[redacted]		[redacted]		[redacted]
[redacted]		[redacted]						
[redacted]		[redacted]						
[redacted]		[redacted]		[redacted]		[redacted]		[redacted]
[redacted]								
[redacted]	[redacted]		[redacted]		[redacted]		[redacted]	
[redacted]	[redacted]		[redacted]		[redacted]		[redacted]	
[redacted]	[redacted]		[redacted]		[redacted]		[redacted]	
[redacted]								
[redacted]	[redacted]		[redacted]		[redacted]		[redacted]	
[redacted]	[redacted]		[redacted]		[redacted]		[redacted]	
[redacted]	[redacted]		[redacted]		[redacted]		[redacted]	

Canterbury's investment in High Processor Count / Low Power HPC will likely see a greater than two-fold increase in sheer numbers of processors available compared to their current BlueGene/L facility, while the architectural improvements in the new system are likely to represent one or more generations of improvement, suggesting significant advances in power efficiency and overall system speed, far beyond the increase in sheet processor numbers.

NIWA and Canterbury will both invest in High Memory / Fast Processor Speed systems, with the NIWA system taking on the demanding production codes (such as climate circulation modelling) that require dedicated and ongoing resources, while Canterbury will support the experimental end of development on this large scale platform, leading to new capabilities and applications.

The annualised investments at Auckland and AgResearch/Otago into Super Clusters are possible due to the smaller increments available within these systems. They can be built up over time, taking advantage in the substantial increases in capacity available through new multi core chip designs, on an annual basis. This has been modelled using a crude measure of a doubling in chip density on an annual basis, based on prior experience in acquisitions at Auckland in recent years.

Currently, the Auckland cluster facility is at about [redacted] of the suggested NeSI investment level in year 1, and is static at this level.

Storage investments will be made on an annual basis with alignment with computational needs. These investments vary in capacities across institutions due to variations in price. Storage typically comes in a variety of specifications, all the way from slow but reliable storage for archiving infrequently accessed data through to highly tuned storage for read intensive or read/write intensive analyses. The costs vary by two orders of magnitude. Storage will be sized appropriately during the tender process.

6.5 The Proposed Operational Budget

The operational cost of NeSI is [redacted] over 4 years, of which the collaborating institutions provide [redacted]. Of the remaining [redacted] is covered by the government investment, and the remaining [redacted] from merit access and cost recovery from support provided to commercial research services to industry. Each of these areas is broken down below, and summarised in Table 7.

Table 7. **Forecast Revenue, Expenditure and Funding** for the Infrastructure (initial 4 years).

[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]					
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]					
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]					
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]

Commercial Services: refers to estimated costs of delivering commercial services.

Management & Governance: includes external Board members, a director and admin staff reflecting the need for strong focused management.

Facilities Administration: Operations & Service team administration at each of four sites to ensure strong alignment and coordination of the highly skilled staff involved in the service teams located with each set of hardware and related services.

Operations: Development and maintenance of layers of software and hardware that together connect to form a coherent, single infrastructure. Enables the best international advances to be adopted for use in New Zealand's research infrastructure, and ensures national interoperability of systems.

Expert Services: Working with domain scientists, specific software tools sourced, developed or deployed for research teams, some of which are generic and applicable to multiple research domains, others require one-on-one service delivery to individual teams.

6.5.1 Revenues

Revenues of [redacted] comprise [redacted] from merit access fees and [redacted] from cost recovery from support provided to commercial research services to industry.

The NeSI partnership expects merit access revenues of [redacted] to be generated through users' research funding. The Board will communicate with research funding and investment agencies and the Ministry of Science and Innovation on the nature of the merit access scheme, whereby only 20% of the costs are required in the first years of the infrastructure. Additional merit access eligible funding sources will be qualified by the Board inside a Merit Access Policy to be established as the Board's first priority on establishment. This allows for appropriate devolved funding mechanisms inside institutions to be qualified as eligible, such as PBRF or CRI Core Purpose Funding-supported research.

A Merit Access pool (equating to [redacted] of resources over 4 years) will be used to support wider sector access while providing necessary training and translation services to ensure researchers can be migrated into the infrastructure effectively. Researchers granted merit access will have [redacted] of related costs covered in this first 4 years, and will pay for the remaining [redacted]

The budget assumes a sliding scale of take-up of the merit pool, acknowledging that the sector will need time to react, and that new capacity will not all be immediately available operationally in any year that it is added. Rates of take-up of merit access used in the budget model here are [redacted]

Commercial revenues will be generated from engagements with the private sector that fall outside of Merit Access. Such opportunities could arise from a variety of sources, from commercial drug discovery to nano-materials simulations. Assessment of the scale of this activity is currently conservative, representing the low levels of high performance computing in New Zealand industry, other than in a few well known cases such as Weta Digital who are already well equipped for their operational needs. Initially, investors will be the providers of resources, with the Board reviewing options to extend commercial relationships to additional collaborators or service providers as needs and opportunities arise. At least the full costs of providing any services inside commercial research contracts will be passed through to any engaged service providers.

6.5.2 Expenditure

The operational costs of NeSI are [redacted] million over 4 years, which includes [redacted] million for management and governance, [redacted] for facilities administration and operation of the distributed eScience platforms, [redacted] million for expert services, [redacted] million for expenses related to support of commercial work and [redacted] million direct operating costs (facilities, power and maintenance) to provide services to the collaborating institutions and the wider sector.

As direct costs related to service provision rise in response to demand, merit access and commercial revenues will be reinvested into scaling up services to meet demand.

6.5.3 Forecast Income and Expenditure Statement

A small operating surplus should be achieved across all years, contingent on the successful implementation of the merit access scheme and on supporting commercially focused research services. The forecast for the first 4-year term is shown at the bottom of Table 7. A full 10 year forecast is available in Appendix 8.

6.5.4 Financial Implications of Operating Principles and Agreements

Section 5 (Governance) provides detailed principles and agreements that give effect to the joint and coordinated form of the infrastructure. Several of these principles and agreements have financial implications.

The partnership's essential role is to operate as a coordination mechanism bringing together capital investment in HPC and eScience nationally. To do so efficiently requires significant investment in the management of teams, systems and software, in a variety of complex technical domains, and across considerable distance. Team resource profiles have been developed that represent a careful approach in terms of the mix of roles specified. This is further discussed in the next sub-section on Operational Staffing Requirements.

6.5.5 Operational Staffing Requirements

An essential principle for the successful coordination of the infrastructure is the establishment of operational transparency, requiring monitoring of utilisation levels across all equipment and within a common framework, supporting audit and accounting functions to manage the Service Level Agreements with each institution. Access to the infrastructure by the wider research sector will be managed through a 'single front door' (i.e. users will access multiple HPC architectures and services through the same common interfaces).

Functional requirements to support the infrastructure include administration of the underlying access policy, the requisite unified access mechanisms integrated into institutional security systems, and support for the various applications and services to deliver capacities to the broad swathe of researchers. Scaling up and broadening scope while refining service levels all requires that operational resourcing needs are carefully considered and provided. The mix of Roles defined within the budget have been established to ensure essential responsibilities are managed and the design goals of NeSI are achievable, all the way from the systems through to engagement with researchers across the NZ research sector.

The allocation of resources at each institution is sized to operate effectively and exploit equipment. Close to whole FTEs will be deployed to ensure efficiency and focus on NeSI services. Every individual will need to have advanced abilities to fulfil a complex set of responsibilities across layers of the infrastructure. In common with other research infrastructures, such individuals are typically highly qualified, either technically or as a researcher, and have significant experience and interest in applying high-performance computing equipment to research problems. The remuneration specified in the budget is commensurate with securing such scarce and necessary personnel.

Table 8 describes the staffing levels proposed, summarised across Institutions and Roles. To coordinate activities nationally with the desired level of integration, a single team approach will be adopted. While specialised equipment will be deployed at various sites, the teams responsible for the maintenance and operation, and those supporting researchers to exploit them, will be combined into a single nationally coordinated team. This ensures efficiency gains through common and shared knowledge bases and practices, with the implication that it will pay to invest in carefully managing a distributed support team.

Table 8. **Staffing** by Institution and Role.

[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]				
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]				
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]				
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]				
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]				
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]				
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]

6.5.6 Investments by Institutions

The Funding Agreement between MSI and the partnership, through the host institution, will require the co-investment outlined in this investment case, to go alongside the government investment. Table 9 provides a full description, broken down by institutions and type of expenditure.

Table 9. **Capital and Operating Expenditure** by Institution.

[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]					
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]					
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]					
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]					
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]					
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]					
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]

The funding agreement will refer to the intent of this Investment Case in terms of capital and operational investments, and will devolve control to implement those investments to NeSI's governance. The partnership, through the host institution, will develop sub-contracts for service level agreements with the collaborating institutions.

The Partnership and Funding Agreements will make explicit the commitments from the institutions in terms of co-invested capital and operating expenditures, which are then combined with the government investment to achieve the levels of investment allocated to each institution.

In arriving at these allocations to institutions a core principle has been to invest in communities where they already exist, to efficiently support and grow existing capabilities. This aligns investments with institutional needs and characteristics, to ensure greatest buy-in and focus, and significantly reduce risks.

6.5.7 Allocations of Resource Capacities

The outcome of the proposed investment is a system of resources coordinated across New Zealand, and made available as capacities allocated to the wider NZ research sector and the investing institutions. The investing institutional allocations recognise the direct investments made by these institutions and their Associate Investors, and provide for commercial access within research contracts. The wider sector allocations support both merit-based research activities, and educational activities.

A summary of these gross allocations is shown in Table 10 below. Note that these figures contain additional income, in the form of co-funding from merit research funding providers, and commercial revenues that generate additional fees. This is the reason for the differences in these figures from earlier tables (where such at-risk revenue was excluded).

In the case where any unused merit access capacity is provided to the commercial sector, NeSI brokers and agrees the terms of the subsequent contract.

Table 10: **Allocations** of Resource Capacities

[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]					
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]					
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]					
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]	[redacted]

7 SUMMARY

New Zealand research stands at a crossroads; the high-speed network is in place, but the remaining pieces of an eScience infrastructure remain elusive. If these can be provided, New Zealand scientists will be on an equal footing to their offshore counterparts, who already benefit from sophisticated research infrastructure. A summary of the rationale for HPC and eScience Infrastructure in New Zealand, from scientific, societal and economic perspectives is as follows:

Science rationale

1. **Keep New Zealand science internationally competitive** in strategic areas by providing the necessary enabling infrastructure. Address the need for New Zealand to be better connected with international science communities.
2. **Researchers able to address previously intractable questions** that require access to HPC resources beyond their current reach. Help researchers directly leverage national investments in HPC as seamlessly as possible, with a minimum of refocusing.
3. **Keep a better record** of our data, experiments, and accumulated knowledge for sharing, validation and for future generations of scientists. Current science records are inadequate for long-term understanding (e.g. in the areas of sustainability, agricultural impacts, and their remediation).
4. **Enable more effective collaboration and resource sharing** between science communities, both within New Zealand and internationally.

Societal rationale

5. **New Zealand researchers trained in emerging eScience practices and do not become further isolated from the global research community**.
6. **Inter-operable products and services** between research agencies enabling better planning and mitigation of pressing issues such as extreme events, climate change and bio-security.
7. **Science as a service**: better access for government departments, educators and general public to science products and related information (e-government).

Economic rationale

8. **Strengthening of scientific outputs** in areas of economic and strategic importance.
9. **Faster research and development times** due to more sophisticated infrastructure, continuously maintained.
10. **Keep New Zealand's top researchers in New Zealand**, **and attract others** to generate innovations and business investment.
11. **Promote efficiencies in research funding** by providing remote access to expensive equipment, virtual meetings, virtual collaborations, shared collections of data and experiments, and a national strategy for creating the support services that researchers require. A delay in developing eScience infrastructure will cost more in the medium and long term.

8 APPENDICES

8.1 APPENDIX 1: National and international case studies

Below are some international comparisons relating to the current investment of other countries in coordinated eScience HPC initiatives. These initiatives have the following in common:

- They stress the need for an end-to-end solution to provide access to high-performance computing,
- They support a sustained, nationally coordinated development strategy and they provide significant resources to engage with specific research communities and help them migrate their work to the new paradigms and facilities that eResearch offers.

Australia

National Collaborative Research Infrastructure Strategy (NCRIS) had in 2008 over $540million AUD to invest in various areas of research infrastructure. The 2009-2013 Australian Budget increased the A$82m portion for eResearch infrastructure (under Platforms for Collaboration) to A$312: A$130m for HPC, A$97m for data storage and collaboration, A$48m to advance the Australian Data Commons, and A$37m for the Australian Research Network.

Canada

Have announced an investment of $88 million for the creation of the first ever pan-Canadian network of high performance computing facilities, a national resource. Seven HPC consortia across the country will be collaborating as full partners in this project.

Japan

Japan is developing a supercomputer (sited at Kobe) that would operate at 10 petaflops, or 10 quadrillion calculations per second. It is sponsored by the government with three main commercial collaborators who plan to complete in 2012 with a development budget of ¥115billion.

Ireland

Has a comprehensive national strategy that coordinates HPC, grid middleware and research communities. Investments from several ongoing and past grants equate to many millions of Euros annually.

USA

Heavy investment since 1990s (TERAGRID), Newly-established Office of Cyberinfrastructure (OCI) at the National Science Foundation (NSF) has 50 currently active projects in 2008. Billions of dollars invested so far, with large amounts budgeted for the future.

United Kingdom

Investment began in 2001 in the UK eScience and eSocial-Science programs, with a first phase budget of £91M, and is now into its second phase. The current UK investment in HPC (HecTOR) is £113M.

S. Korea

A nationally-coordinated programme in eScience and HPC is currently underway, targeted at 5 specific national science priorities.

8.1.1 Examples of Successful HPC Centres

The following examples describe successful overseas HPC Centres that utilise some of the models described above:

CSC Finland

The CSC operates as a wholly government owned non-profit company. It is administered by the Ministry of Education (Department for Education and Science Policy). CSC's statutory administrative bodies consist of the General Meeting of Shareholders (AGM), the Board of Directors and the Managing Director. The operative management of the company is under the responsibility of the Managing Director together with the Management Group.

The duties of the CSC's Board of Directors are defined through the Finnish Limited Liability Companies Act. CSC has no committees. The activities of the Management Group include planning the company's strategic policies, creating annual operational plans including the budget, and planning investment schemes. It also allocates resources, decides on key activities and significant operational issues, and monitors the implementation of the decisions made. Scientific advisers support the Managing Director's and Management Group's work. Their tasks include assessing and developing services provided by CSC, representing academia in the planning process of CSC, as well as enhancing CSC's contacts with the national and international research environment.

This operating model has notable elements including:

- Research service focus: research is supported across different phases, including modelling the problem, selection of suitable methods, software and databases, across a range of disciplines
- Forum: various national and international forums are supported by CSC, including computational science, high-performance computing, IT infrastructures, networking, and security. In addition, CSC participates in various scientific and computational forums
- Service development: ongoing development of services offered by CSC is achieved through cooperation with customers, and tertiary institutions, research institutes and the business community.

Victorian Partnership for Advanced Computing (VPAC)

VPAC operates as a non-profit registered research agency, established in 2000 by a consortium of Victorian Universities. It provides expert services, training and support in Advanced Computing as well as professional R&D services in the application of Advanced Computing in the fields of Engineering, Geophysics, Health, Life Sciences, Spatial Information and Grid Computing. VPAC also maintains collaborative relationships with R&D Organisations, Government Department and other Associations, and has a wide range of partners and clients.

The governance model operated by VPAC includes a Director, Board and Committees. The Board meets quarterly to set strategy and review the company's operations in accordance with the company's Constitution, and is advised by an Audit & Risk Management Committee on financial reporting, accounting controls, risk, audit, regulatory and internal controls compliance etc. Other committees include Member Support Committee and an Operations Committee, with each committee consisting of a representative from each Member University appointed by the Board of Directors. The Member committee deals with member funding and guidelines, suggests how VPAC can support member activities, and participates in the review process for HPC purchases.

The Operations committee reviews operations of the VPAC Supercomputing Systems, reviews usage and distribution, represent users and presents their suggestions, complaints and problems, works with VPAC to develop projects, and authorises creation of new VPAC projects. Access for University members is defined by a usage quota based system to ensure users get sufficient capacity, which is determined by the number of users per University.

8.2 APPENDIX 2. Background on NeSI's Principal Partners

8.2.1 BlueFern HPC Facility

BlueFern®, a collegial, HPC and eResearch services facility based at the University of Canterbury, is a key component of the current NZ HPC and eScience environment. It has been in existence since mid-2006, beginning with an IBM p575 based computer system and augmented in mid-2007 by the addition of the first IBM BlueGene in the southern hemisphere. Currently, the BlueFern infrastructure consists of two IBM supercomputers, a P-Series and a BlueGene, housed in a secure, reliable data centre at the University of Canterbury. The facility was included in the Top100 supercomputer list in 2008.

BlueFern set out to provide modern scientific computing capability and associated skills to New Zealand researchers and their partners. It acts as a national facility through both BeSTGRID and direct accessibility via the KAREN network. BlueFern's capacity is of a scale that enables internationally competitive research. However, within three months of implementation and go-live, both the BlueGene and other systems were fully utilised.

BlueFern has provided the basis for a national HPC centre to support advanced research activities in New Zealand based on best practices of national HPC centres located in Europe, North America and Australia. It is managed on a daily basis by a director and two front-line support staff focused on helping and supporting researchers using the facility.

BlueFern foundation partners include:

- University of Canterbury
- The MacDiarmid Institute
- Victoria University
- Auckland University of Technology

8.2.2 BeSTGRID eScience and Grid Computing Capabilities

Led by the University of Auckland, BeSTGRID is New Zealand's national HPC grid service, providing the grid infrastructure for New Zealand research and science. It provides eScience infrastructure in a shared, distributed resources model for research support not only for national projects but also for also for international collaborations.

BeSTGRID established this shared infrastructure through hosting large-scale resources and implementing a variety of middleware along with computational and data-intensive science applications and services. The services offered are focused on top-down science drivers (mainly Biosciences and Geosciences), and bottom-up service development and capability building.

Initial funding was provided by the TEC Innovation and Development Fund, from late 2006 through to early 2008. The TEC contract was led by Auckland University, with subcontracts held by Massey and Canterbury universities. BeSTGRID is currently funded out of the MoRST eResearch Vote RS&T allocation. Funding runs out towards the end of 2010.

BeSTGRID coordinates shared eScience infrastructure across a growing partner network that includes:

- AgResearch
- Auckland University
- AUT
- Canterbury University
- IRL (MacDiarmid)
- Landcare
- Lincoln University

- Massey University
- Otago University
- Victoria University
- Waikato University

Through support of member institutions' Offices of Research Management, IT Directorates, and in some cases dedicated Centres of eScience, BeSTGRID runs as a collaborative eScience infrastructure network with strong ongoing participation from member institutions. BeSTGRID has in place experienced, skilled, software development and support, and managerial staff at key NZ research institutions, along with strong collaborative leadership from domain scientists.

As a fully functional eScience community, BeSTGRID is actively engaged in supporting the development of eScience infrastructure by providing national leadership in establishing the necessary middleware, processes, and skills-base to support shared resources and research collaborations. It has mature, inclusive governance and management structures, which include two science advisory committees, a steering committee, and a technical working group, all with well-established systems of collaboration and communication.

8.2.3 NIWA High Performance Computing Facility

NIWA has had HPC capability since 1999, for use in water and atmospheric modelling to support public good research in New Zealand and international programmes such as IPCC. Due to NIWA's time-critical needs, it maintains an ownership model based on facility management at NIWA.

NIWA has utilised their HPC to support the major New Zealand capability in the following research and applications:

- Improved weather, hazard and environmental forecasting
- Climate modelling
- Ocean modelling
- Atmospheric chemistry modelling
- Fisheries modelling

NIWA has recently invested in an upgrade of their HPC infrastructure and storage; the new supercomputing facilities are located at NIWA's Wellington site at Greta Point. The NIWA HPCF (High Performance Computing Facility) will be accessible to the wider research community under the following principles:

1. The HPCF is a major strategic asset for NIWA and NZ, much like the research vessel *Tangaroa*. Its primary purpose is to enable NZ scientists to carry out nationally and internationally significant research and run time-critical production workloads that require access to a capability-class supercomputer, and to enhance international research collaborations.
2. HPCF user information services, including documentation, real-time system information and issue tracking, will be seamlessly accessible to NIWA and External users.
3. User requirements will inform management of the HPCF.
4. HPCF Projects will contribute towards the running costs of the HPCF on a fair and equitable basis.

The Facility is overseen by the HPCF manager, with 24/7 support staff and scientific programming assistance as part of the core service. It is overseen by a governance group Chaired by the Chief Scientist (Atmosphere) and includes the GM Information Technology and other relevant Chief Scientists. With appropriate storage and communications, the facility has been scoped to meet demand for 5 years with a mid-life upgrade to double the capacity. Access will be through the infrastructure proposed here.

8.2.4 AgResearch Grid

AgResearch has been a leader in high performance computing in New Zealand with its initial foray into the field in 2000 leading to its inclusion in the Top500 supercomputer list in 2003. Through this period the AgResearch facility and its staff have maintained AgResearch's position at the leading edge of molecular and bioinformatics based research, including large-scale microarray design, manufacture, experimentation and analysis, and proteomics analysis. More recently, from 2006 to the present, this facility was the main resource underpinning assembly of the first drafts of the sheep, deer and fungal endophyte genomes, and will over the next five years lead the data processing and bioinformatics required for assembly of further pastoral forage plant species, and increasingly, individual genomes of some thousands of genetically valuable livestock.

This facility has been critical in establishing and maintaining a number of national and international collaborations over this time, starting back in 2001 with the Victorian Department of Primary Industry in Melbourne for whom AgResearch provided computational and bioinformatics services and consultancy; Nutrigenomics New Zealand, a collaboration between the University of Auckland and two other CRI's, for whom AgResearch has developed and host database, compute and collaboration services; and similar compute, database and collaboration services for Sheep Genomics Australia (a collaboration between around 10 Australian Universities and State and Federal Government Sheep research programs); Plant Genomics Consortium, supporting clover and ryegrass bioinformatics and molecular research; Via Lactia (Fonterra); The International Sheep Genomics Consortium, the international collaboration that oversaw the first draft of the sheep genome and the development of a commercial 60K SNP chip. More recently AgResearch, in collaboration with BeSTGRID and with the support of a REANNZ capability build fund, has developed the concept and implementation of www.biocommons.org.nz, a collaboration resource specifically targeted and supporting inter-institutional and international collaboration on the analysis and interpretation of the results of Next Generation Sequencing experiments.

Currently AgResearch's main high throughput compute cluster consists of around 150 cores used for a mix of database hosting, parallel bioinformatics and statistical data pre-processing, and downstream bioinformatics and statistical based data-mining of DNA and protein sequences. The computational facility is currently in the process of a substantial upgrade in capacity in the area of medium to large memory (500Gbyte), multiprocessor (48 core) SMP machines to meet the requirements of the latest breed of graph based genome assemblers, and large statistical models in support of Genome Wide Selection design and production analyses. As part of the NZGL initiative AgResearch is proposed to be the key provider of IT resources to support a wide range of bioinformatics analysis tools. It is envisaged that a significant Supercluster facility and associated network storage will be provisioned to support genomics activities across New Zealand as part of these NZGL developments.

8.3 APPENDIX 3: MMRF-Green Model

The MMRF-Green model is a multi-regional, dynamic CGE (Computable General Equilibrium) economic impact model. It distinguishes up to eight Australian regions (six States and two Territories) and, depending on the application, up to 144 commodities/industries. The model recognises:

- domestic producers classified by industry and domestic region;
- investors similarly classified;
- up to eight region-specific household sectors;
- an aggregate foreign purchaser of the domestic economy's exports;
- flows of greenhouse gas emissions and energy usage by fuel and user;
- up to eight state and territory governments; and
- the federal government.

The model contains explicit representations of intra-regional, inter-regional and international trade flows based on regional input-output data developed by Monash University's Centre of Policy Studies, and includes detailed data on state and Federal government budgets, with each region modelled as a mini-economy. Second round effects are captured via the model's input-output linkages and account for economy-wide and international constraints.

The MMRF-Green model estimates the implications of policy (and resource) changes on an inter-industry and year-by-year basis. For the purposes of this exercise MMRF-Green can track economic impacts in terms of:

- national output, including breaking down the results to identify activity level changes for retailers, domestic industries associated with the manufacture of lightweight bags, re-usable bags, bin liners and replacement bags, and other upstream and downstream industries;
- employment (that is, changes in full time equivalent jobs);
- net impact on consumers (that is, by measuring the change in real net welfare);
- government budget positions;
- real wage levels; and
- the balance of trade (that is, regional international export earnings, international import expenditures and international balance of payments).

The model also provides a detailed representation of the energy sector and associated resource flows. In particular, it provides insight to the usage patterns of coal, oil and gas across the economy (the latter being important to plastics production) and forestry activity (important for paper production). The model also estimates the greenhouse gas implications of different production outcomes.

8.4 APPENDIX 4: Needs Analysis. Results from the national survey

Sample results from survey questions follow. A full set of results is obtainable on request.

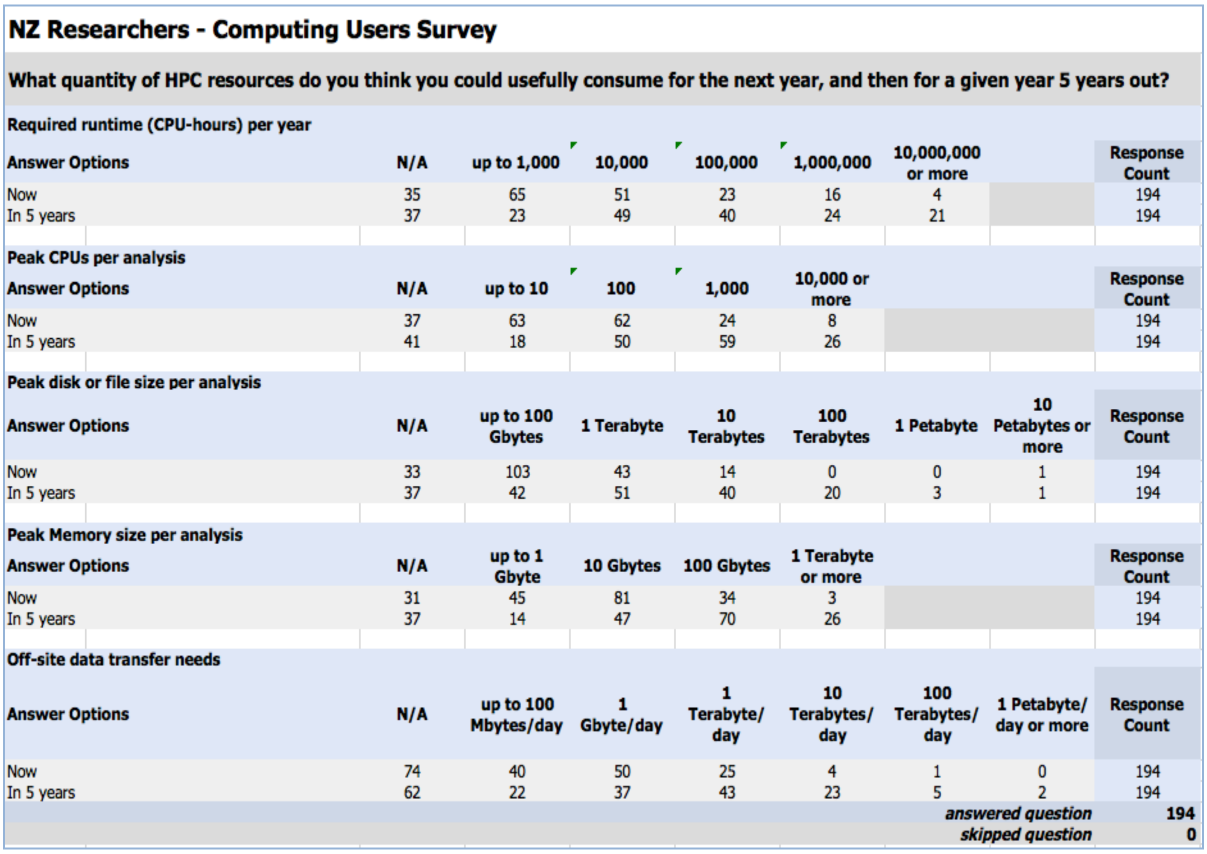

NZ Researchers - Computing Users Survey

What quantity of HPC resources do you think you could usefully consume for the next year, and then for a given year 5 years out?

Required runtime (CPU-hours) per year

Answer Options	N/A	up to 1,000	10,000	100,000	1,000,000	10,000,000 or more		Response Count
Now	35	65	51	23	16	4		194
In 5 years	37	23	49	40	24	21		194

Peak CPUs per analysis

Answer Options	N/A	up to 10	100	1,000	10,000 or more			Response Count
Now	37	63	62	24	8			194
In 5 years	41	18	50	59	26			194

Peak disk or file size per analysis

Answer Options	N/A	up to 100 Gbytes	1 Terabyte	10 Terabytes	100 Terabytes	1 Petabyte	10 Petabytes or more	Response Count
Now	33	103	43	14	0	0	1	194
In 5 years	37	42	51	40	20	3	1	194

Peak Memory size per analysis

Answer Options	N/A	up to 1 Gbyte	10 Gbytes	100 Gbytes	1 Terabyte or more			Response Count
Now	31	45	81	34	3			194
In 5 years	37	14	47	70	26			194

Off-site data transfer needs

Answer Options	N/A	up to 100 Mbytes/day	1 Gbyte/day	1 Terabyte/ day	10 Terabytes/ day	100 Terabytes/ day	1 Petabyte/ day or more	Response Count
Now	74	40	50	25	4	1	0	194
In 5 years	62	22	37	43	23	5	2	194
							answered question	194
							skipped question	0

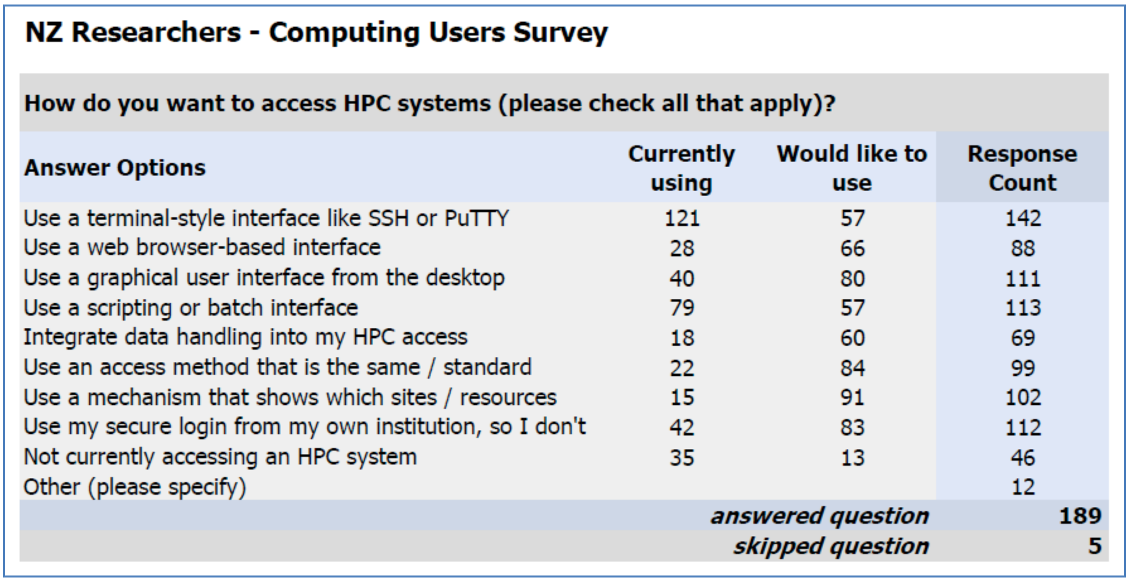

NZ Researchers - Computing Users Survey

How do you want to access HPC systems (please check all that apply)?

Answer Options	Currently using	Would like to use	Response Count
Use a terminal-style interface like SSH or PuTTY	121	57	142
Use a web browser-based interface	28	66	88
Use a graphical user interface from the desktop	40	80	111
Use a scripting or batch interface	79	57	113
Integrate data handling into my HPC access	18	60	69
Use an access method that is the same / standard	22	84	99
Use a mechanism that shows which sites / resources	15	91	102
Use my secure login from my own institution, so I don't	42	83	112
Not currently accessing an HPC system	35	13	46
Other (please specify)			12
		answered question	189
		skipped question	5

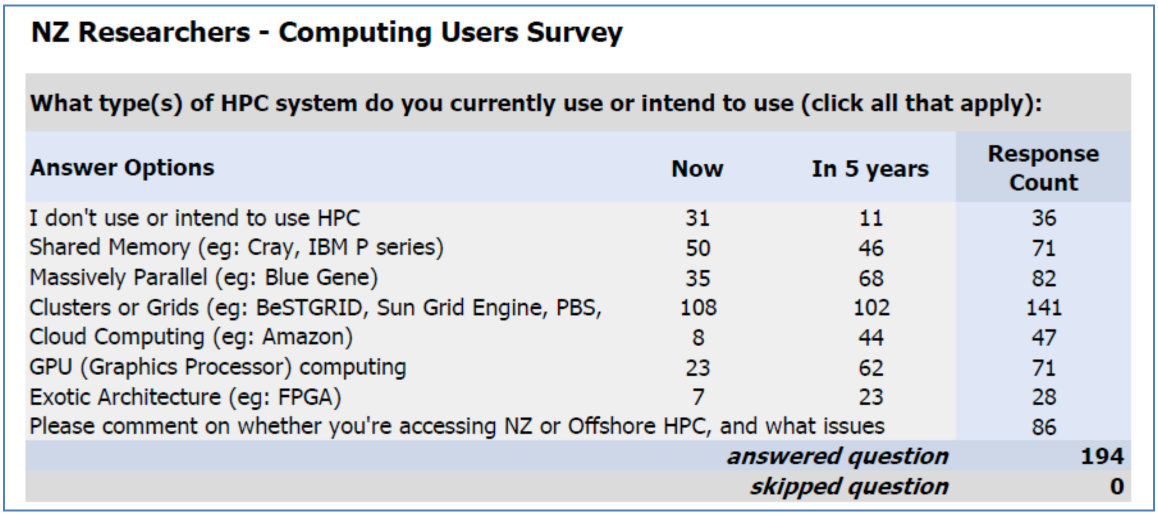

NZ Researchers - Computing Users Survey

What type(s) of HPC system do you currently use or intend to use (click all that apply):

Answer Options	**Now**	**In 5 years**	**Response Count**
I don't use or intend to use HPC	31	11	36
Shared Memory (eg: Cray, IBM P series)	50	46	71
Massively Parallel (eg: Blue Gene)	35	68	82
Clusters or Grids (eg: BeSTGRID, Sun Grid Engine, PBS,	108	102	141
Cloud Computing (eg: Amazon)	8	44	47
GPU (Graphics Processor) computing	23	62	71
Exotic Architecture (eg: FPGA)	7	23	28
Please comment on whether you're accessing NZ or Offshore HPC, and what issues			86
		answered question	**194**
		skipped question	**0**

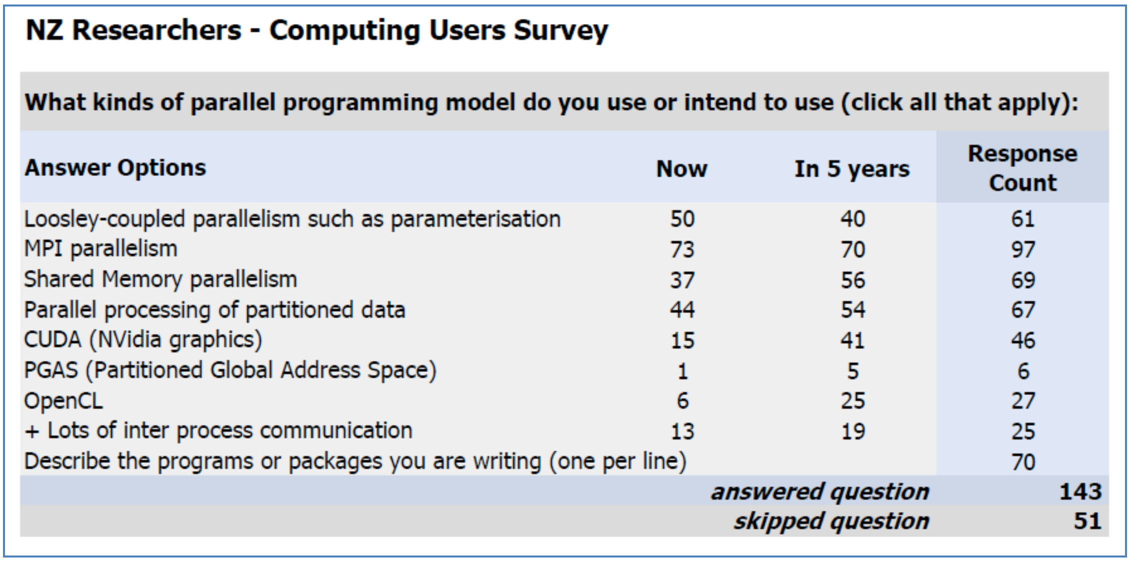

NZ Researchers - Computing Users Survey

What kinds of parallel programming model do you use or intend to use (click all that apply):

Answer Options	**Now**	**In 5 years**	**Response Count**
Loosley-coupled parallelism such as parameterisation	50	40	61
MPI parallelism	73	70	97
Shared Memory parallelism	37	56	69
Parallel processing of partitioned data	44	54	67
CUDA (NVidia graphics)	15	41	46
PGAS (Partitioned Global Address Space)	1	5	6
OpenCL	6	25	27
+ Lots of inter process communication	13	19	25
Describe the programs or packages you are writing (one per line)			70
		answered question	**143**
		skipped question	**51**

8.5 APPENDIX 5: Foundation Charter

NeSI has drafted a Foundation Charter (Charter), consisting of a *Statement of Purpose*, a set of agreed *Objectives*, *Key Success Factors*, *Role of the Board*, *Communications*, and a set of *Rights and Obligations*. These elements of the Charter should be read collectively in order to gain a full understanding of how NeSI intends to function. They inform and provide the basis for the MoU, Funding Agreement and Service Level Agreements.

8.5.1 Statement of Purpose

The core business of the National eScience Infrastructure is to coordinate capital investment in HPC and eScience efficiently and operate it effectively for access by the NZ research sector.

8.5.2 Objectives

The following is a summary of the objectives of NeSI, thus will form the major priorities for the Board's governance to tackle.

1. **Enhance Sector Capabilities**: Build a well-connected research sector supported by significantly increased and efficiently delivered HPC and eResearch capacities and delivered by a strongly governed and coordinated infrastructure.
2. **Coordinate HPC and eScience Assets**: Efficiently and effectively coordinate and support the management and operation of significant HPC and eResearch assets.
3. **Provide Fair and Simple Access**: To lower the barrier to accessing sophisticated HPC and eResearch infrastructure, while providing transparency around policy and priority setting.
4. **Be research partners with HPC and eResearch expertise**: Be an organisation respected as a model of collaborative, distributed team-work, known for openness and transparency in relations and operations, and efficient and effective as research collaborations.

Further discussion on these Objectives, within each category, are below, followed by exposition of related Rights and Obligations.

Enhance Sector Capabilities:

- To develop, foster and operationalise a shared national HPC vision, focused on co-ordination between partners, with strong governance. The vision is of a collaborative approach to HPC and eResearch provision nationally, with evidence of changes in behaviour around research infrastructure investments.
- To significantly grow NZ HPC and eResearch capacities to closely match the needs of the NZ research community. Principal Investors will own resources, and investments will include heterogeneous platforms matched to needs as assessed periodically.
- To create a fair and transparent Access Policy that allows for any researcher at a NZ research institution to apply for access to the range of resources and services within the infrastructure. Equity of access shall be determined by policy and oversight from the NeSI Governance Board.
- To efficiently use HPC and eResearch resources, through advising and training researchers so they can identify and access the most suitable platform(s) that match resource capabilities to the research problem at hand.
- To provide high-quality HPC and eResearch learning materials and training opportunities to all members of the NZ research sector, through commitment to pursue complementary and collaborative education and training activities.

Coordinate HPC and eResearch Assets:

- Offset the high risks and liabilities that would otherwise accrue to investing institutions, enabling each partner to commit to ongoing investment in resources and service provision, by guaranteeing purchase of agreed resource capacities from each NeSI partner. All equipment will be fully owned or leased by a host institution.

- Host institutions to negotiate with the NeSI partnership on suitability of platforms offered. NeSI to govern how funds for future upgrades are disbursed based on projected need and uptake—for at least the portion of the platform funded by NeSI.
- Teams and services will be integrated nationally and centrally coordinated, with operational transparency of monitoring and utilisation of all resources to support Service Level Agreements with each institution.

Provide Fair and Simple Access:

- NeSI to provide a range of resource pools for established researchers, emerging researchers and research targeted education and outreach.
- To achieve fair and simple 'single front door' access through developing an Access Policy for the infrastructure, describing access to resources via federated institutional identity credentials and through tools, applications, and services that meet researcher needs.
- The Access Policy aims to define a set of queues representing investors and education, publicly funded and commercial research communities, and through these queues manage access to the computational platforms to meet all the various research sector needs and priorities while managing the diversity of these requirements.
- The NeSI Board is responsible for determining a fair and equitable access strategy, which will ultimately set policy (the Board may set up a separate committee for this and any other task as they deem necessary). The NeSI Board will also ensure that major NeSI resources are not captured by specific institutions or science projects.

Be research partners with HPC and eResearch expertise:

- Develop, grow and sustain a nationally distributed team dedicated to the delivery of HPC and eResearch services.
- Attract and develop a team with abilities and desire to engage in research collaborations as a contributing research partner.
- For the team to consist of dedicated (1.0 FTEs by default) people to ensure efficiency and clarity of purpose. The intention of having co-investment coordinated within the infrastructure is to retain the full capacity of the team coordinated in common directions, to achieve focus, and ultimately effectiveness as a team.
- Be transparent in strategy, policy, investments, and operations, reporting to the wider sector on plans and activities openly, to develop trust and reduce opportunities for capture or monopolisation of resources by vested interests.

8.5.3 Key Success Factors

The following are factors will describe the success of this venture of the medium to long term. The Board will develop strategies, policies, and plans what reflect on these factors and design them in to the core of the organisation.

1. **Reflective, Strong, Lean Governance**: A tightly knit and focused governance Board is considered essential to the partnership, with explicit well defined and agreed shared objectives and governing principles. This Charter addresses this factor directly.
2. **Sophisticated Team Development**: The people constituting NeSI will face an ambitious set of challenges, from working as a distributed team, through coordinating across very advanced information technologies, to developing in-depth working knowledge of and relations with a broad range of research problems. Being successful in this environment requires considerable talent, self-awareness, and breadth of experiences beyond both traditional science and information technology.
3. **Considered and Consultative Engagement**: NeSI will have working relations with researchers of varied backgrounds, funding agents with specific priorities, policy analysts, and the administrations of institutions across the sector. Success will require engagement that

considers the complex forces within the sector and compromises that deliver on the varied stakeholders needs.

4. **Productivity Focus**: Each of these communities have their own definition of productivity, whether it be research productivity in terms of papers or discoveries, capabilities developed within the sector, or expertise applied with an industry focus and specific tangible commercial outcomes or economic gains. Focus on productivity is essential to NeSI's success.

8.5.4 Scope of Activity

The scope of activity is initially defined by the following constraints:

1. **Geography, Sector**: New Zealand and its research sector, with a primary focus on publicly funded research, and a secondary focus on commercial research.
2. **Resource Domains**: High Performance Computing, specialised distributed computing, data intensive research, and related resources to support a fully formed HPC and eScience infrastructure.
3. **Financial Scale**: Financial constraints limit the scale of the investment, and hence while NZ research sector members are eligible to apply for access, the size of the community being supported at any specific time will be limited.

8.5.5 Role of the Board

The Charter forms the conceptual basis on which all investors agree to govern NeSI, and provides the basis on which the Board will balance the interests of its stakeholders, including its lead investing institutions, government agencies, the wider New Zealand research sector, and especially the communities of researchers that will come to depend on it. The Board is responsible for achieving the objectives, strategies, and policies, and the Charter will be used to shape strategy and all investment and policy development, and guide and inform operational decisions.

The Board will: (1) develop NeSI's strategic plan, (2) monitor operations and development activities relative to the strategic plan along with annual and quarterly objectives, (3) review effectiveness of the Management Group's strategies, (4) approve and monitor large projects, (5) shape corporate governance policies and practices, and (6) monitor compliance with these policies and practices.

The strategic plan should cover a period up to a minimum of 3 years and include the Board's thoughts on the key objectives of NeSI over that period, including how it will manage growth and operationalise investments, and which communities it will prioritise allocation of resources to, and include expected staffing requirements and capital allocations.

Further details on the structure and responsibilities of the Board are detailed in Section 5.2.1.

Board Committees

The Board may establish committees with specific responsibilities as and when needed. Likely examples are (i) to ensure fair and equitable access to NeSI services and (ii) to plan for ongoing and future investments in computational systems.

8.5.6 Communications

The intention is for NeSI to present a united face to the research sector, to be identifiable as the primary contact point for HPC and eResearch in New Zealand. This requires the investors to commit to a tightly aligned approach to branding, communications, and publicity and promotions.

The Board will ensure that management has appropriate policies in place to facilitate effective communication processes, satisfy principles and objectives for unity, transparency and consultation, and ensure that financial results and other material events are reported on a timely basis.

8.5.7 Rights and Obligations

The incentives and obligations described below have been developed by NeSI's Transition Working Group as operating principles for the infrastructure. Note that each pair is coupled—a right carries an obligation, sometimes two. Together they provide a framework to align partners and investors with the overall Goals and Charter of NeSI, and thereby ensuring that no single investor or partner captures disproportion value or monopoly control over NeSI or its invested in resources and capabilities.

To Achieve Strategic Intent:

Intent	Description	Right / Obligation
Coordinate: International	Strengthened connections to international research	Right
Coordinate: International	Connect with external / international communities Leverage best of breed methodologies and technologies	Obligation
Transparency: Operations	Operational reporting/monitoring is common and open	Obligation
Transparency: Use of Capabilities and Services	Reassurance of equality of access Reinforces maturity through enabling scrutiny and feedback	Right
Promote: Common brand (NZ Inc approach)	Commit to a common brand for NeSI activities. Shared promotion and outreach programmes	Obligation
Promote: Attractiveness and Sophistication of Sector	Coherent services and capabilities Broader scope of applications and services on offer Increased attractiveness of research sector	Right
Govern: Participate	Consult and be transparent in relevant policy formation and planning Proactive participant in governance	Obligation
Govern: Balance and Validation	A coherent and consulted sector wide approach	Obligation

To Provide Appropriate HPC Assets:

Intent	Description	Right / Obligation
Procure: Coordinate	Shared planning for procurement Assurance of fit for purpose Transparency of contracts and agreements	Obligation
Procure: Economies of Scale	Enables scaling up to new research challenges Provide coherent and diverse HPC capabilities and services Greater buying power Neutrality in purchase process	Right
Invest: Shared capabilities and services	Co-invest institutional resources into infrastructure Exploit sector wide investments, not just local investments Comply with agreed standards for identity	Obligation

	management and middleware	
Invest: Return on investment	Leverage of combined investment from sector and government Managed financial risk for large investments Stronger incentives to exploit investments across sector, leading to better uptake of services, platforms, and facilities	Right
Sustain: Reinvest	On an annually reviewed basis, commit to reinvestment through procurement planning process On an annually reviewed basis, commit to Service Level Agreements around investments	Obligation
Sustain: Ongoing Investment	Adapt platforms procured to meet emerging needs	Right

To Deliver HPC Infrastructure Access:

Intent	Description	Right / Obligation
Coordinate: Access Mechanisms	Managed access available, on differentiated rates to be agreed, to all NZ public & private sector researchers Common access mechanisms Comply with agreed standards for identity management and middleware	Obligation
Coordinate: Access	Ease of access to all resources Ease of use Lower barriers of entry / activation	Right

To Deploy Human Resources:

Intent	Description	Right / Obligation
Develop: Human Capability	Development of research and infrastructure capabilities Exposure to leading edge of research paradigms Training next generation of researchers	Right
Develop: Human Capability	Shared education and outreach programme	Obligation
Mature: Capabilities, Services	Larger teams with broader skills Sustainable services Managed risk Less fragmentation of skill bases and teams	Right
Coordinate: Capabilities	Matrix management of teams Shared planning, development, and delivery	Obligation

8.6 APPENDIX 6: Recommendations on Priorities

In preparing this Investment Case, many issues have arisen and been dealt with in a manner that establishes the principles and basis for formation of the partnership and establishment of NeSI. However there are several outstanding issues that need attention, and that are appropriate to defer to the Board for resolution during the establishment phase of NeSI. The key issues identified by the Transitional Working Group are noted and commented on below, and passed on to the Board for their consideration.

8.6.1 Definition and Prioritisation of Merit Access Scheme

Within the Merit Access Scheme, and initial definition of the qualifying research communities has been approached, in the broadest of senses. Within this definition will need to sit a prioritisation, as it is expected that over the course of the life of NeSI demand for resources from the wider sector will outstrip the capacities available. It remains to firm up this definition and ensure it meets with the overall objectives of this proposal and the sector.

8.6.2 Coordination with Funding Providers

To achieve the sector contributions to the Merit Access Scheme, in the form of Merit Fees revenue, a possible efficiency of approach is to coordinate with funding providers. An initial conversation has been held with the Chair of the Marsden Fund, who sees this as a desirable strategy that removes duplication and reduces transaction costs to the sector. It remains to establish whether this approach is viable, and if so how it proceeds.

8.6.3 Coordination with Researchers and institutions

Another approach to improving sector take-up of merit access requires working with individual researchers, research teams and institutions early in the process of bidding for research funding, so that NeSI services can be costed appropriately into the budget. Such activity might be coordinated in many ways:

1. Informal direct contacts with researchers.
2. A formal process for researchers to engage with NeSI.
3. Collaboration between NeSI and the research offices and commercialisation offices at the various research institutions.