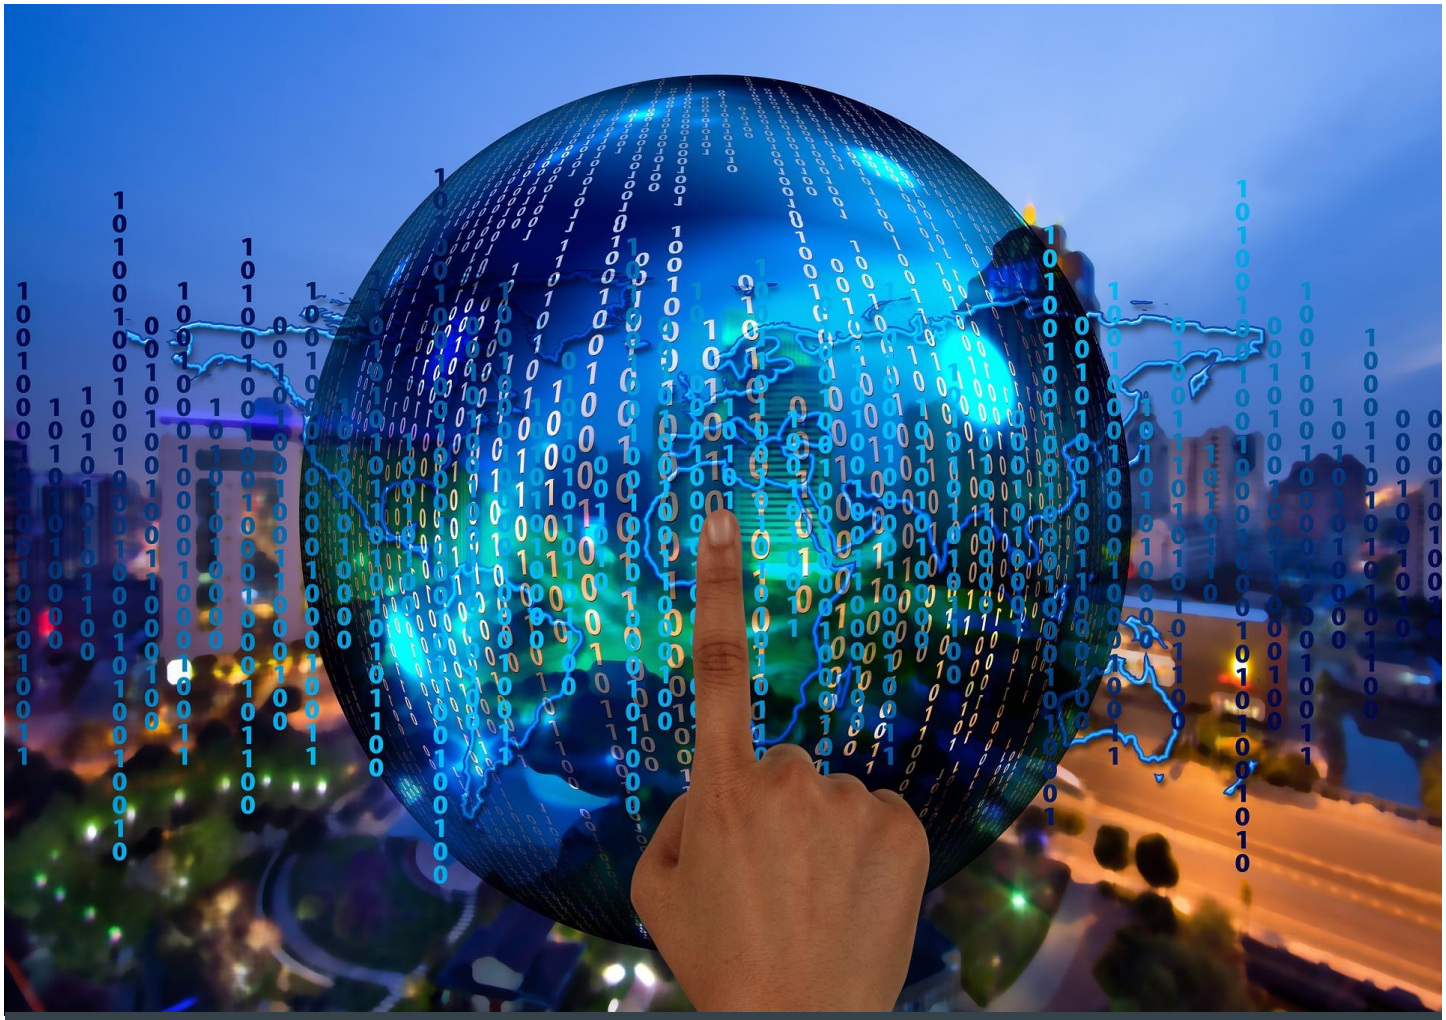

# Evaluation of the New Zealand eScience Infrastructure (NeSI)

## Summary Report

**11 August 2017**

ALLEN+CLARKE

# 1. BACKGROUND

NeSI was established in 2011 providing a nationally coordinated high-performance computing (HPC) network to enhance the capability and quality of New Zealand's research. NeSI's establishment was in response to an undersupply of supercomputing facilities in New Zealand and was intended to overcome coordination failure in New Zealand's science system. The provision of centralised infrastructure intended to provide a platform for greater research collaboration both domestically and internationally, and to increase the competitiveness of New Zealand's research overseas.

NeSI has a mandate of supporting research outputs which are key to New Zealand's economic prosperity and the public good. It is expected that NeSI has sound alignment with government science priorities, including the National Statement of Science Investment and the Strategic Science Investment Fund, and that NeSI increases the scope of research being undertaken.

NeSI was established through an initial Crown Funding Agreement between the Crown and the principal investors of the Universities of Auckland and Canterbury, and the Crown Research Institutes, NIWA and AgResearch. AgResearch withdrew from the principal investors in 2011 followed by increased involvement of Landcare Research and the University of Otago.

The second phase of NeSI (known as NeSI2) commenced in 2014 combining Crown funding ($27 million to June 2018) and funding from the principal investors ($26 million). The University of Canterbury were a party to the second phase of NeSI, but due to exceptional circumstances following the earthquakes withdrew from the collaboration in 2016.

During 2017 NeSI, in partnership with NIWA, has contracted to purchase replacement HPC equipment (NewHPC) to start operating in early 2018. From the preliminary information available, NewHPC will be considerably more powerful than its predecessors and will be sited in a primary data centre at NIWA with a disaster recovery site at the University of Auckland. Unlike the original 2011 configuration, NewHPC includes substantial data storage capacity and data analysis services.

# 2. PURPOSE, OBJECTIVES AND METHODS OF THE EVALUATION

The Ministry of Business, Innovation and Employment (MBIE) is seeking transparency as to what its investment in NeSI has achieved, and independent advice to inform its future funding decisions in HPC infrastructure. Ahead of any funding decisions, MBIE is seeking:

- assurance and transparency as to what investment in NeSI has achieved; and
- independent advice to inform future funding decisions in national HPC infrastructure, capacity, support and development.

Specific objectives of the evaluation are to assess:

- the **effectiveness** of NeSI, and its **impact** on research capability, capacity and outcomes (including the benefits for both research and researchers);
- the **relevance** of NeSI's services, including its alignment with the government's strategic science investment priorities and the way that NeSI has shaped these priorities (including the role that NeSI has played in supporting New Zealand's international science engagement in terms of both research and infrastructure investment collaboration);

- the **sustainability** of NeSI, including its ability to adapt to changing external circumstances;
- the **value for money** achieved through the government and collaborating partners' investment in NeSI; and
- the **appropriateness** of NeSI's business structure and operations.

A range of methods was used to evaluate NeSI, including:

- **review of key documents**
- **analysis of NeSI data** related to HPC capacity and usage.
- **interviews** with a range of stakeholders.
- **interviews** with 21 NeSI users and 8 potential users.

The collected data was analysed thematically through the use on NVivo software and was reviewed against the evaluation framework to support value judgements on NeSI's effectiveness, impact, relevance, sustainability, value for money, and, appropriateness.

# 3. FINDINGS AND CONCLUSIONS

## 3.1. Effectiveness

NeSI is mostly delivering effective services, which are used by a range of research communities. Most of the researchers engaged in this evaluation who are using NeSI stated that NeSI was an effective service that meets their HPC needs. NeSI's user support services are seen as valuable, providing assistance to researchers on the computational aspects of using NeSI, and helping to run applications. These services are effective for helping new users to gain a basic level of HPC competence, and for assisting established HPC researchers to run their compute effectively. There is a need to increase the extent of use across research institutions, with use currently concentrated in a limited number of institutions.

### How effectively has NeSI delivered what was intended?

NeSI is providing well-regarded and effective HPC and support services. The number of individual users is increasing, and these researchers come from a range of research disciplines. The researchers using NeSI are drawn from a limited number of institutions. Until 2016, the bulk of NeSI core hour usage was consumed by the largest investor institutions: University of Auckland, NIWA, and the University of Canterbury. There has been little use by the other universities, although Massey and Victoria Universities' use has recently increased following NeSI's implementation of training and awareness raising activities intended to broaden sector uptake. CRI use of NeSI is lower than use of similar facilities by comparable entities in Australia.

NeSI is investing in a new HPC platform that will substantially increase the amount of capacity available. However, the use of NeSI by a limited number of institutions and individuals suggests that there is a need to stimulate demand. NeSI is aware of this, and is delivering well-regarded training workshops that are helping to increase researcher computing capability. These are particularly effective for teaching a basic level the HPC skills of early career researchers, such as post-graduate students, as well as established researchers new to HPC.

NeSI's user support services are seen as valuable, providing effective assistance to researchers on the computational aspects of using NeSI, helping to run applications. These services support researcher productivity, allowing them to focus on their research rather than technical computing issues. The support services are effective for helping new users to gain a basic level of HPC competence, and for assisting established HPC researchers to run their compute effectively.

There does appear to be a gap in NeSI's service offering in terms of assisting researchers to use NeSI to enhance the quality of their science output by offering specialist support to research disciplines that have specific requirements for HPC. This will be particularly important as the volume of research data available to disciplines such as genomics and bioinformatics is predicted to increase substantially.

## 3.2. Impact

Measuring NeSI's impact is difficult due to the considerable time lag between the research being undertaken and its impact. Attribution of any impacts to research infrastructure is even more challenging. The evaluation therefore focused on identifying changes in the behaviour of researchers that is likely to lead to impacts in the longer term. Researchers are reporting productivity gains, increased participation in national and international collaboration, and global recognition for NeSI-related research.

### What impact has NeSI had on New Zealand's research sector?

The evaluation found that NeSI has contributed to positive outcomes for researchers and research institutions. The availability of a national HPC infrastructure has allowed research institutions to attract and retain high quality scientists.

It has allowed researchers to conceive of and investigate larger scale problems, many of which are too big to be addressed by an individual researcher or small teams of researchers. This is promoting a culture of collaboration. NeSI provides the technology to allow New Zealand researchers to participate in multi-national consortia.

The evaluation identified examples of reported researcher productivity gains. Compared to their previous computing arrangements, researchers reported that NeSI is enhancing:

- the **speed** of research. The average speed increase reported was that research could be undertaken six times faster on NeSI than on researchers' previous computing arrangements;
- the **scale** of research, particularly modelling and simulations in climate change, earthquake simulations, chemistry and physics. NeSI enables these researchers to investigate larger problems of greater complexity through undertaking longer modelling/simulation runs and more runs;
- the **accuracy** of research, particularly when moving from a PC or a small local cluster to NeSI.

NeSI-related research has been recognised at a global level, helping to enhance New Zealand's international standing in science. Researchers reported publishing impactful papers in highly ranked journals.

## 3.3. Relevance of NeSI

NeSI demonstrates relevance as a response to the HPC needs of the research sector in New Zealand. Its functionality and services meet researcher requirements and compare favourably to alternative options for computer-assisted research. It has processes in place to identify and respond to changing needs.

The government has set out its priorities for the New Zealand science system in the National Statement of Science Investment (NSSI), which lists five specific goals. There is evidence that NeSI is assisting the science community to deliver on these goals. It supports:

- a **better-performing science system** (NSSI goal 1) by increasing research productivity, enabling participation in scaled-up research with greater potential impact, and providing training to create a pipeline of researchers that have the capacity to undertake high quality HPC-related research;
- **growth in business enterprise research and development spending (BERD)** (NSSI goal 2) through targeting commercial entities with research and development functions to invest in NeSI under the subscriber model;
- **reduced complexity and increased transparency** (NSSI goal 3) by offering a consolidated HPC service, replacing individual entity-owned and managed HPC infrastructure;
- **continuous improvement in New Zealand's international standing as a high-quality R&D destination** (NSSI goal 4) through being used in world-leading research papers and the supporting researcher ability to participate in multi-national collaborations; and
- **evaluation and monitoring of performance** (NSSI goal 5) though its publicly-reported KPI data and current evaluation of its performance.

### To what extent is NeSI relevant to government and research sector priorities and needs?

NeSI is assisting the research community to deliver on the goals set out in the National Statement of Science Investment. There is alignment between the outputs of NeSI-related research and government priorities in conservation, climate change and Maori development.

However, NeSI's alignment with government priorities has been relatively ad hoc. If there is an expectation that NeSI will be a tool for delivering on government priorities it would be beneficial to take a more strategic approach.

NeSI is offering a more relevant service than the available alternative options for most HPC-competent researchers. It is seen as offering more timely and accessible support than overseas HPC facilities; superior accuracy, speed and scale to local clusters; and better affordability than cloud services. The availability of high quality support services is a key 'value add' that NeSI offers that the alterative options cannot match.

NeSI will need to consider how it remains relevant to the research community as technological advances change research needs. Emerging areas of need are how to store, manage and transport data. Under its platform refresh, NeSI's partners have purchased hardware with the capacity to provide cloud computing, data storage and data management services, and to support the introduction of virtual laboratories. However, implementing these services would represent a

substantial increase in NeSI's offering and would need to have appropriate supports to enable researchers to use them effectively. Any service expansion will need to be carefully planned and resourced appropriately.

## 3.4. Sustainability

NeSI is reliant on Crown funding to remain viable and would be unlikely to survive should government investment be discontinued. International experience suggests that it is unlikely that any research HPC service will be able to become completely independent from some form of government investment. NeSI has demonstrated that it can survive shocks, successfully adapting to the withdrawal of AgResearch and the University of Canterbury as principal investors. Its platform refresh means it is now well-placed to grow its user base and become more sustainable. To do so, it will need to address weaknesses in its relationship with sector partners and ensure a reliable high bandwidth network service remains available.

### To what extent is NeSI, and its achievements, likely to be sustainable?

NeSI relies heavily on Crown funding, which provides it with a secure base income which it can use to attract further investment. The portion of Crown-funded portion of NeSI's income has not reduced since its inception, and the limited revenue-generating opportunities in a research environment mean it is unlikely to move towards becoming a self-sustaining entity in the near future.

NeSI has proven to be a resilient organisation in that it was able to survive the withdrawal of AgResearch and the University of Canterbury, which resulted in the loss of 5 FTE from its research support service. However, NeSI has not been successful in attracting additional collaborators during its six years of operation and its sustainability would be compromised if another of the current collaborator organisations withdrew.

NeSI has been less successful at developing relationships with key partner organisations. While NeSI has attempted to coordinate its services with REANNZ and NZGL, there were underlying tensions between the organisations and some competitive behaviour and overlaps in service provision.

NeSI is operating in a dynamic environment, with rapid changes to technology and evolving research needs. To date, NeSI has responded appropriately to the availability of changing technology. It employed a range of consultation mechanisms to seek sector input into its platform refresh, and its new infrastructure platform to be launched in 2018 has the capacity to deliver new service lines that will be sought by researchers over the next few years.

## 3.5. Value for money

While it is too early to definitively determine the value achieved from investment in NeSI, there are encouraging signs that NeSI is likely to offer net benefits over time.

### To what extent has the total investment provided value for money?

Cumulative expenditure from government and collaborator institutions on NeSI was around $70 million from 2011 to 2016, with a further $21.5 million expenditure projected for 2017. There are signs that the value provided by NeSI will outweigh its costs.

Over the time that NeSI has been operational, there has been a trend of lowered costs and a growing user base. The 'price' of NeSI, as measured by unit costs, will reduce when the new NeSI platform is launched in 2018, with efficiencies achieved through a national procurement, reduced number of data centres and advances in technology results in an increase in the number of cores that can be purchased per dollar. NeSI's unit costs are on a par with Australian benchmarks.

While not directly applicable to New Zealand, which is at an earlier stage of HPC usage maturity, studies in the US and EU show substantial returns for each dollar of investments in HPC. Applying the study ratios to New Zealand suggests a return of $390 - $690 million for every $10 million invested. Although this is likely to be an over-statement of the total net benefit, it does indicate that investment in HPC has been demonstrated to offer value overseas and is likely to do so in New Zealand.

Overall, although evidence is limited, we conclude that NeSI is performing well against proxy indicators of value, and its performance is improving over time.

## 3.6. Appropriateness of NeSI's business model

NeSI's business model allows the collaborator entities to have a degree of control over their investment and ensures that it remains focused on its research purpose. The NeSI allocation model is not functioning effectively, with the merit allocation being over-subscribed in the most recent three quarters, and the subscriber allocation significantly under-utilised. The availability of subsided merit access to NeSI acts as a disincentive to becoming a subscriber or collaborator. Some stakeholders expressed a view that certain entities were 'gaming the system' by using a large volume of core hours under the merit access category and, in a moral sense, should be paying for it.

### To what extent are the business model, structure and operating approach fit for purpose?

NeSI's legal status as an unincorporated joint venture is more appropriate than alternative legal structures. The unincorporated collaborative model allows for investing entities to set NeSI strategy, helps ensure it remains focused on its research purpose, and is more resilient to market changes than a commercial entity. There are few clear drivers to move to an alternative structure and NeSI will benefit from a period of stability after the recent disruption caused by managing the University of Canterbury withdrawal, and to allow the changes made to its business model during the transition to NeSI2 to bed down.

NeSI's business model has not been successful in attracting new investors, with one entity withdrawing and no new collaborators being brought in over NeSI1 and NeSI2. The subscriber model has not yet performed as intended, its allocation substantially underutilised.

There is an absence of clearly defined incentives to invest in NeSI; the business model allows for researchers to access NeSI through the merit allocation, which is oversubscribed. This provides little motivation for entities to pay for a service that their researchers can access for free, other than the moral obligation to be a 'good citizen'. The balance of incentives will need to be addressed if NeSI is to attract more investors as collaborators or subscribers.

## 4. RECOMMENDATIONS

The evaluation makes the following recommendations:

1. That NeSI retain its current legal structure but enhance its governance by:
    a. Increasing diversity of the Board; and
    b. Focusing on setting NeSI's longer term strategic direction.
2. That NeSI lead the development of a strategy, supported by MBIE, that identifies needs, gaps, and targeted and prioritised responses (including service offerings) in relation to:
    a. NeSI's role in offering cloud computing, data storage and access, and data movement services, including consideration of whether it is best suited as a provider or broker, or whether other entities are best suited to deliver the service;
    b. NeSI's role in offering necessary informatics, applications, tools and data expert and virtual laboratory support for targeted research communities such as genomics and the environment, itself or in support of or in partnership with other organisations; and
    c. More explicitly aligning its service to supporting the research community to deliver on government science strategies, priorities and programmes.

    The strategy should clearly define NeSI's role within the eScience eco-system, and consider appropriate service delivery models and resourcing needs to fulfil this role.
3. That NeSI develop strategies to attract more investment, including how it can better incentivise entities to become collaborators and subscribers. This is needed to support NeSI's growth and sustainability. Options for consideration include:
    a. Limiting the merit allocation to 20 percent of total capacity, and implementing a competitive process for the allocation;
    b. Limit NeSI use for post-graduate students and proposal development for non-collaborator entities to the subscriber allocation category; and
    c. Continuing to seek usage from commercial entities engaged in research as subscribers.
4. That NeSI focus on expanding the breadth of users (including institutions and researchers within these institutions) by:

a. Identifying which research entities in the science system could benefit from improved access to HPC (e.g. which CRIs, which CoREs and NSCs), and developing a targeted strategy to engage those that are not currently using NeSI;

b. Continuing to deliver awareness raising activities such as supporting the ResBaz events and holding NeSI promotional events at collaborator or user entities, and

c. Improving its tracking of user categories and research projects to monitor use across relevant government strategies and priorities and science system platforms such as CoREs and NSCs.

5. That NeSI strengthen its support services by:

    a. Further promoting the consultation service, provided by the NeSI Computational Science team; and

    b. Considering, as part of future strategic planning, the benefits of building specialist science capability within the support service in key research discipline areas, such as genomics, bioinformatics, chemistry and physics.

6. That NeSI grow stronger partnerships with:

    a. Relevant science system platforms including CoREs, Regional Research Institutes, and NSCs; and

    b. MBIE-funded strategic science infrastructure investments.