NeSI
New Zealand eScience Infrastructure

2017

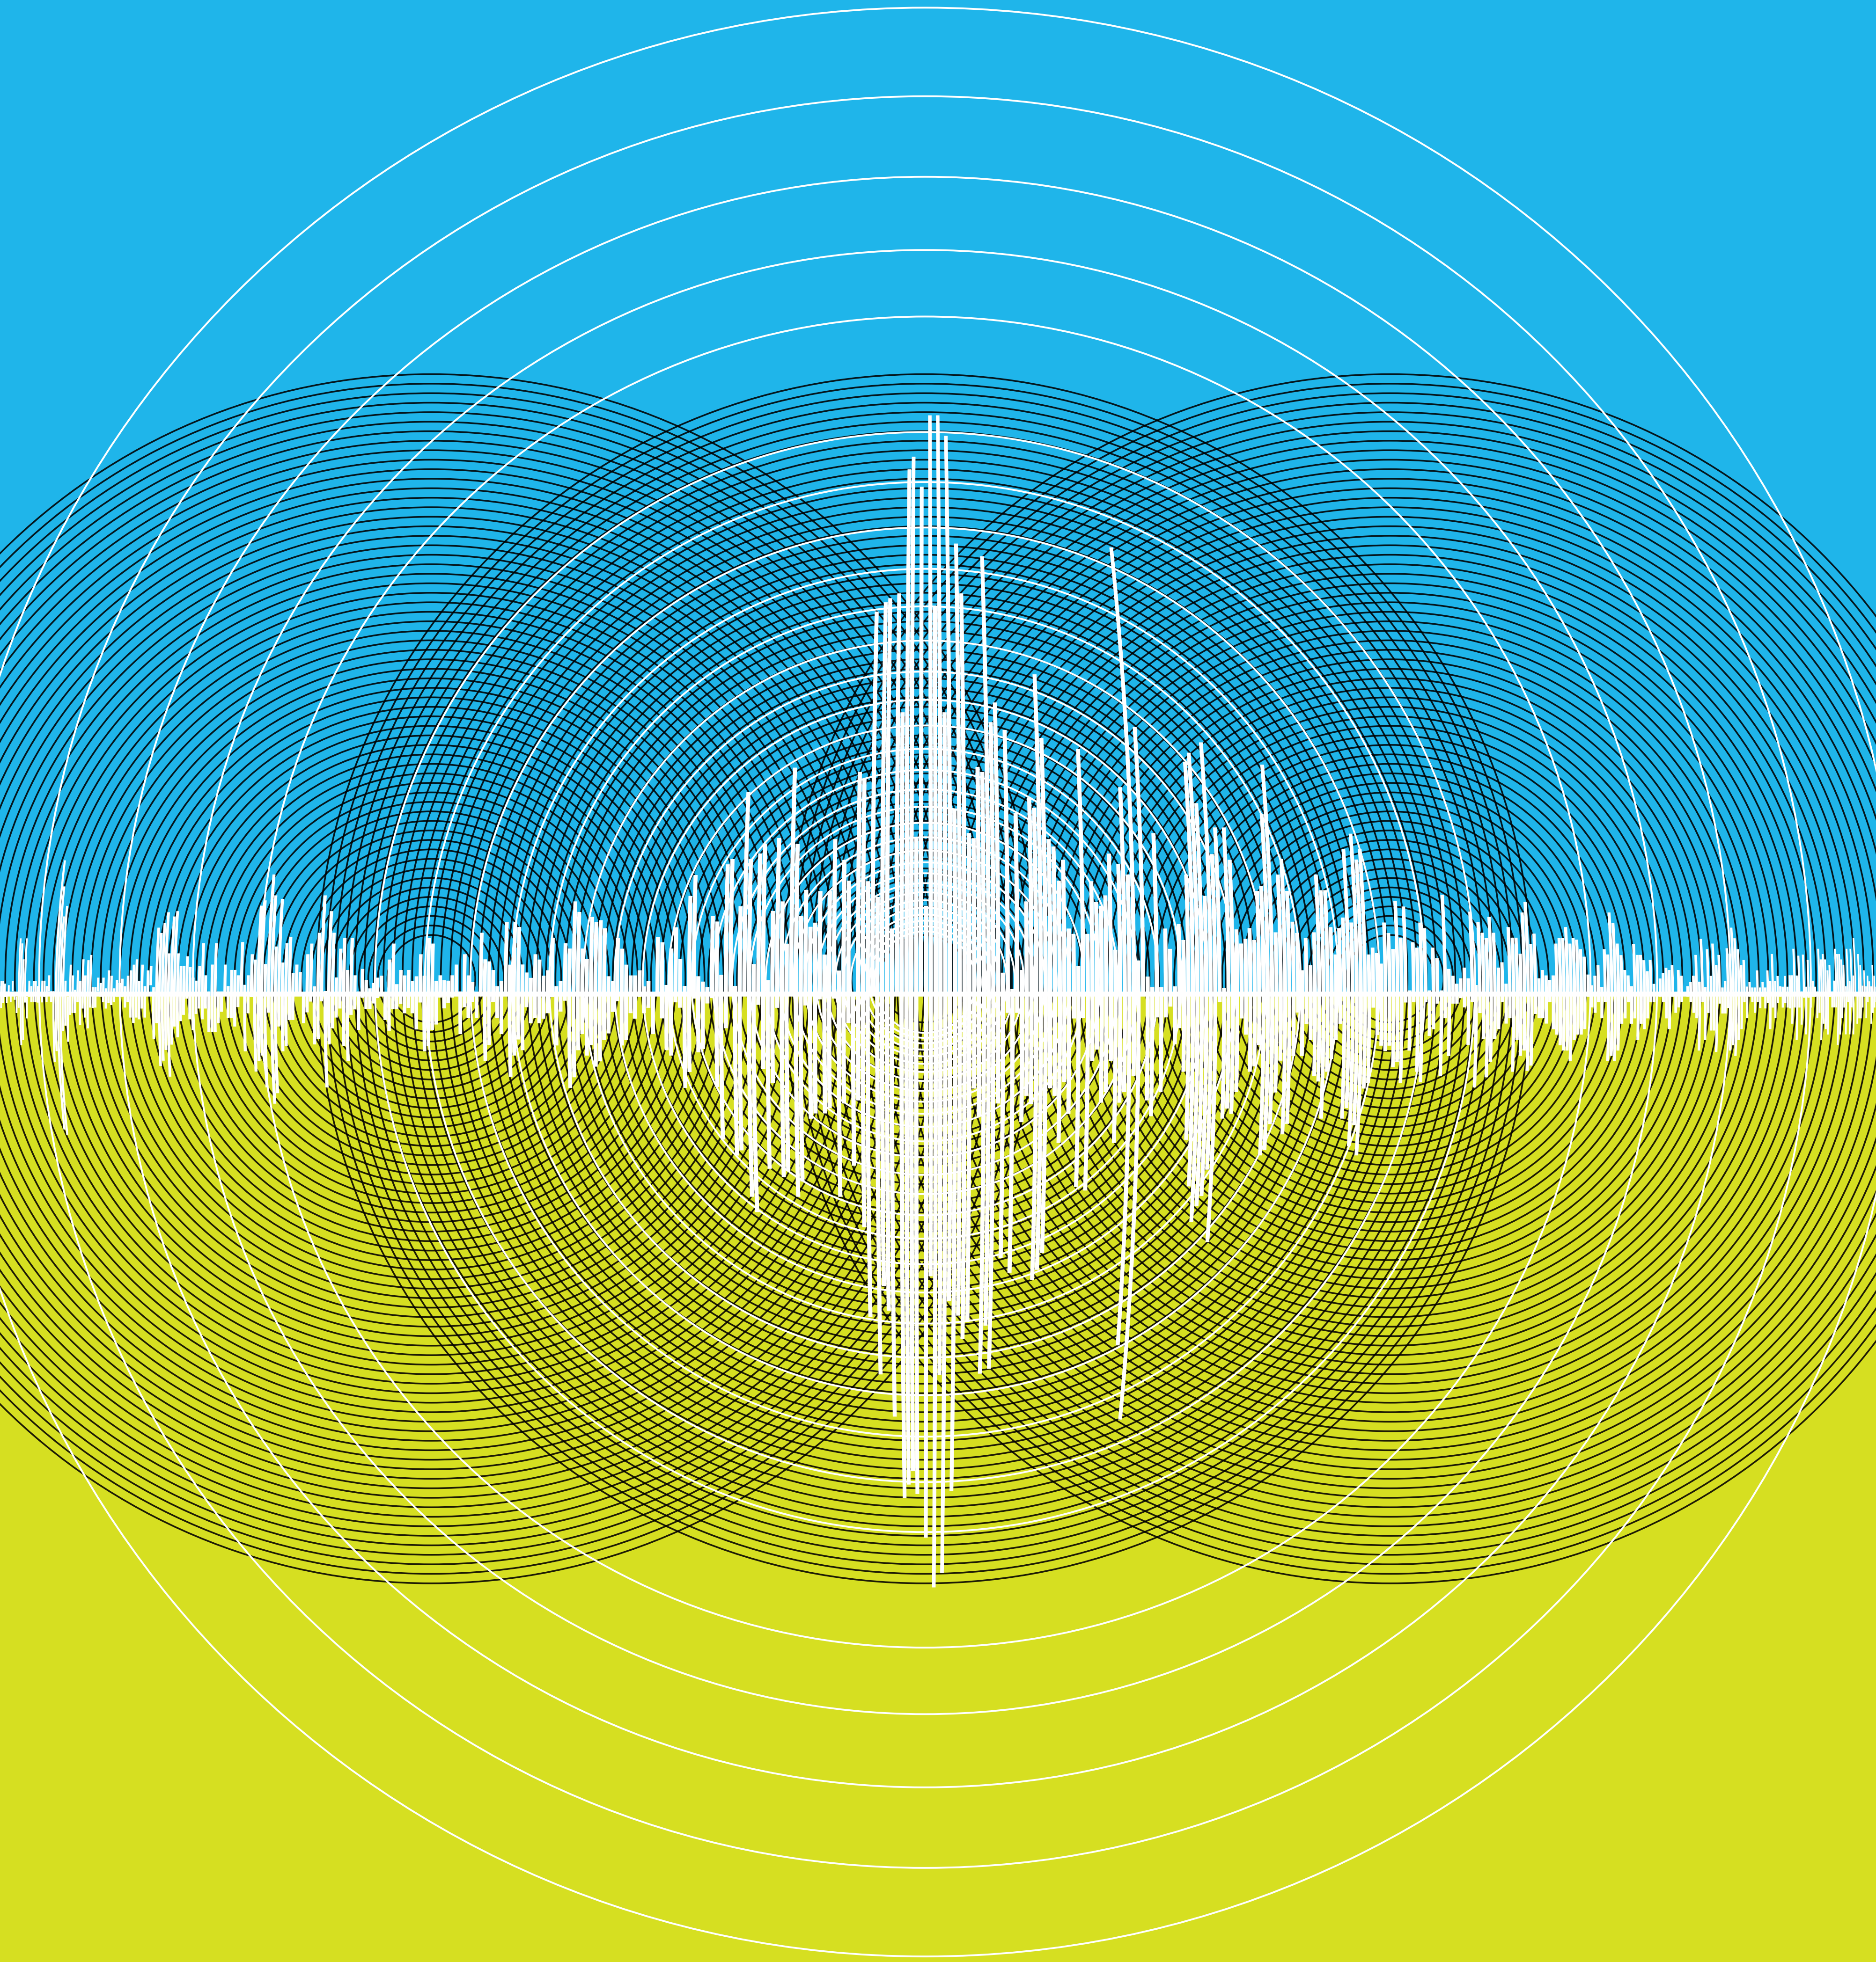

Our Advanced Digital Capabilities

Objectives

NeSI
New Zealand eScience Infrastructure

2017 has been yet another milestone year for NeSI, marking our first six years of existence.

Looking back, we have advanced hugely from our initial fragmented and geographically devolved network of three nodes with an eclectic mix of collaborators and users, through to a shared asset ownership model with vastly increased capacity and diverse capabilities. We have an enhanced collaborator support model and an increasing inventory of successful case studies from a wide range of research communities. We are shortly to commission a very advanced and consolidated national facility delivering significant increases in performance.

At the same time, we have trained a large number of individual researchers and laid the groundwork for wider rollout of Software Carpentry courses and an improved understanding of the potential for High Performance Computing (HPC) across more research areas than ever. There is a new emphasis on data, its collection, storage, curation, and analysis which is becoming ubiquitous across many fields of research in New Zealand.

During 2017, MBIE completed a comprehensive evaluation of NeSI's first six years from which we emerged with considerable credit, but also received some pointers for future changes to guide our next six years. We will be addressing these, in some detail, over the next few months as we develop the Business Case for NeSI 3 and a new contract to commence in mid-2019.

Meanwhile, the NZ Science System has continued to evolve and will be more demanding than ever on a viable eScience infrastructure, not least to analyse and process the avalanche of data that drives much of our discovery and applied science. NeSI is clearly a major player in this, and we will continue to work and consult with our researchers, collaborators and subscribers, and strengthen our relationships with allies such as REANNZ, Genomics Aotearoa, and MBIE to provide our science communities with the services they need to succeed.

In conclusion, I must thank my board and management for their efforts through what has been an intense but rewarding period in NeSI's history. Next year promises to be at least as busy though I sense we are ready for whatever it brings, strategically and operationally.

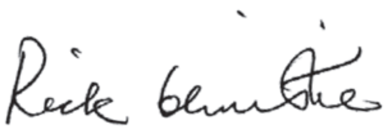

Rick Christie
Chair, Board of Directors

"There is a new emphasis on data, its collection, storage, curation, and analysis which is becoming ubiquitous across many fields of research in New Zealand."

DIRECTOR'S REPORT

Looking back on 2017, there was a strong sense of NeSI hitting its stride. The team balanced a demanding infrastructure procurement programme, supported a comprehensive external evaluation of our progress and impact across our six-year life to date, saw significant growth in the scale of research and researchers supported, all while improving on already high levels of user satisfaction.

As 2017 closes, our infrastructure plans are turning into reality. Our interim platform is in place at the Tamaki Data Centre, allowing the entire platform operations to shift there while capital works are carried out to prepare the Greta Point High Performance Computing Facility for its new role. In early 2018, it will assume duty as the primary site for New Zealand's most advanced computing platform, a significant milestone for the sector, and a major achievement for NeSI, its collaborators and partners.

Across the middle of the year, NeSI worked in support of the external panel tasked with evaluating NeSI's first six years. The panel made recommendations in six areas, covering governance, longer term strategy, scope of services, breadth of uptake, partnerships, and sustainability. Overall, the panel indicated NeSI had met expectations in most objectives, and their comments on areas for future work have been well received. NeSI has an action plan agreed with MBIE, and we're on the journey to framing our future.

A highlight for me this year has been the opportunity to work more closely with team members and researchers. As we refresh our infrastructure, we're also surfacing new and improved services. Doing so is reopening discussions with researchers on their needs and expectations of NeSI, with these voices becoming a driving force for the directions we take. As we commission our new infrastructure early in 2018, I'm personally looking forward to our collective shift in focus; we'll have an eye to exploiting a uniquely powerful advanced computing platform, and an eye to the many ambitious goals of New Zealand's researchers. Underpinning the government agenda for research, we're well positioned to support new directions and bold aspirations across a high performing research system.

Nick Jones
Director

> "Underpinning the government agenda for research, we're well positioned to support new directions and bold aspirations across a high performing research system".

Progress against our strategic goals.

MBIE Evaluation Recommendations

During 2017, MBIE carried out an evaluation of NeSI's first six years. A summary report from the evaluation is publicly available, and the recommendations are listed below.

Governance:

1. Increase diversity of the Board.
2. Set NeSI's longer term strategic direction.

Develop a strategy, supported by MBIE, that identifies needs, gaps, and targeted and prioritised responses (including service offerings) in relation to:

1. 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.
2. 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.
3. Aligning NeSI's services to supporting the research community to deliver on government science strategies, priorities and programmes.
4. Defining NeSI's role within the eScience ecosystem, and consider appropriate service delivery models and resourcing needs to fulfil this role.

Develop strategies to attract more investment, including how it can better incentivise entities to become collaborators and subscribers. Options for consideration include:

1. Limiting the merit allocation to 20 percent of total capacity, and implementing a competitive process for the allocation.
2. Limiting NeSI use for post-graduate students and proposal development for non-collaborator entities to the subscriber allocation category.
3. Continuing to seek usage from commercial entities engaged in research as subscribers.

Focus on expanding the breadth of users (including institutions and researchers within these institutions) by:

1. Identifying which research entities in the science system could benefit from improved access to HPC (e.g. which CRIs, CoREs, and NSCs), and developing a targeted strategy to engage those that are not currently using NeSI.
2. Continuing to deliver awareness-raising activities such as supporting the ResBaz events and holding NeSI promotional events at collaborator or user entities.
3. 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.

Strengthen support services by:

1. Further promoting the consultation service, provided by the NeSI Computational Science team.
2. 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.

Grow stronger partnerships with:

1. Relevant science system platforms including CoREs, Regional Research Institutes, and NSCs.
2. MBIE-funded strategic science infrastructure investments.

NeSI has worked with MBIE to prioritise its action plan for responding to these recommendations, with some addressed during the remainder of 2017, and the majority forming a core component of the work plan from 2018.

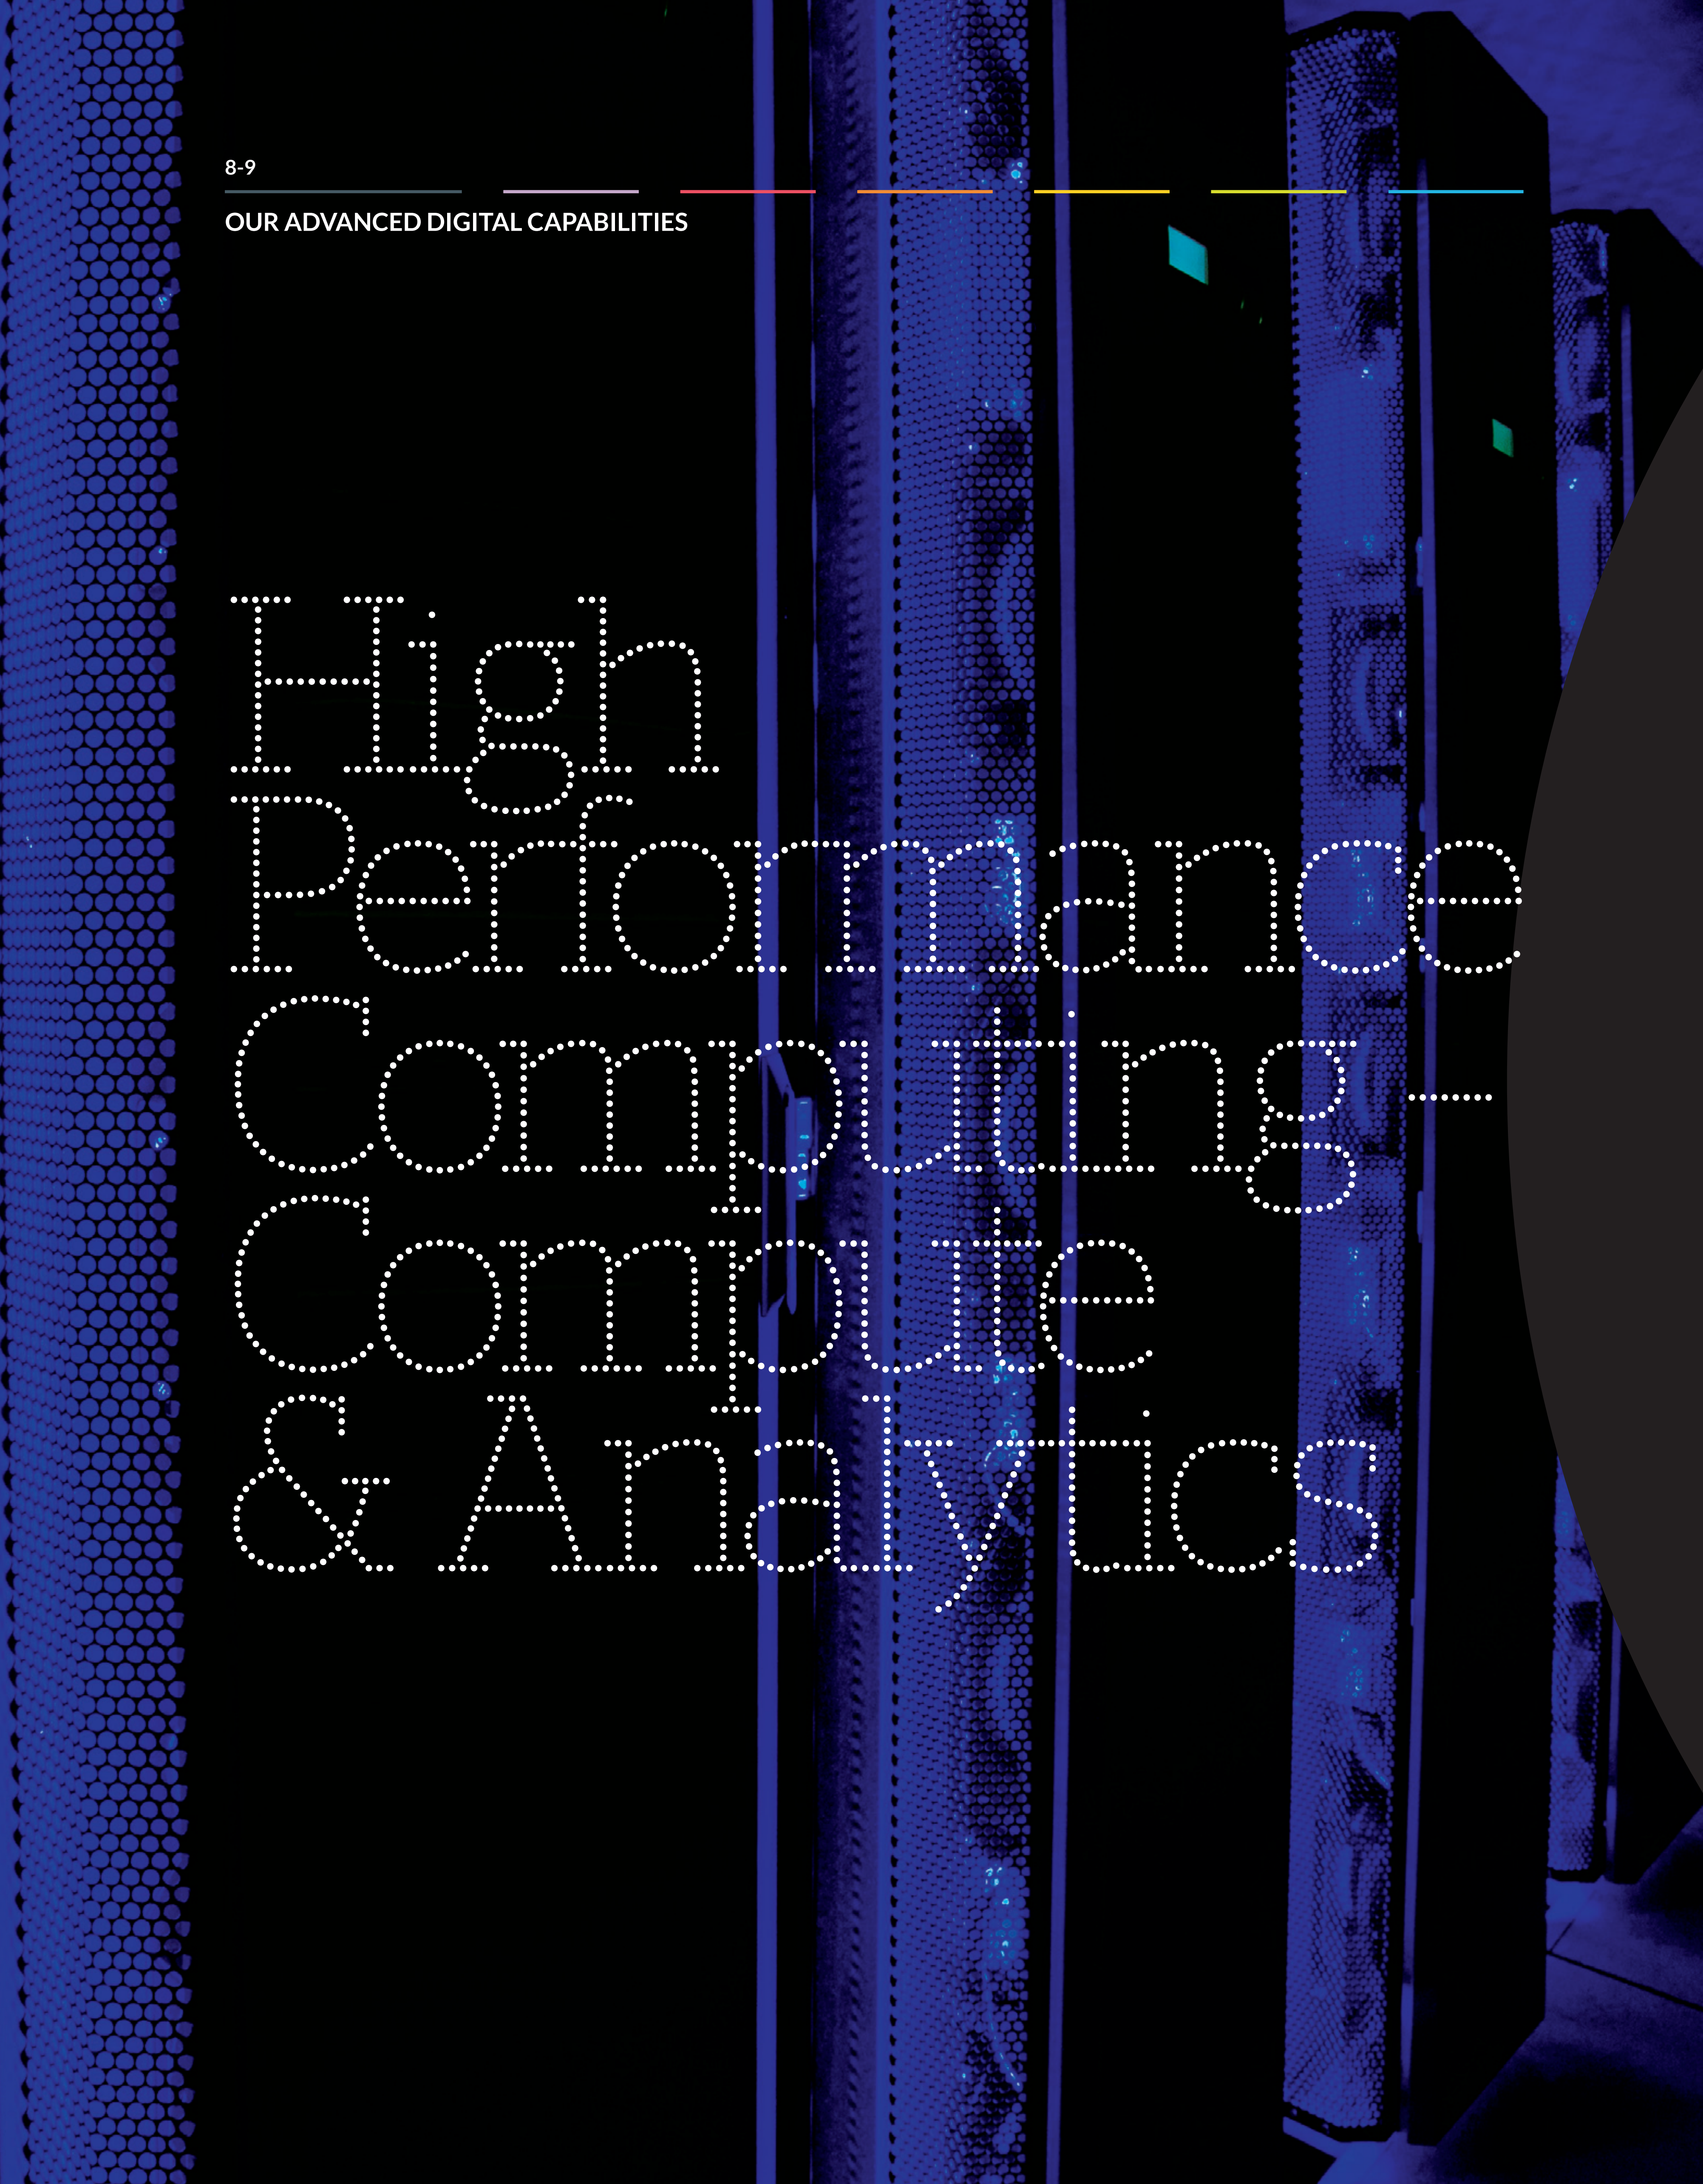

High Performance Computing – Compute & Analytics

Driven by the needs of researchers for high-performance productivity.
NeSI integrates advanced digital capabilities into a range of eScience services, and ensures advanced computational research projects are backed by the power necessary to make them a reality.

High performance computing – compute and analytics.
NeSI makes high performance computing (HPC) available to New Zealand researchers, lifting their ability to tackle large, complex problems, and enabling research to be carried out much faster. NeSI's facilities are the country's most powerful computers for scientific computing.

Development of HPC capabilities during 2017.

As 2017 drew to a close, NeSI celebrated the completion of some major milestones in the refresh of its platforms. Cray was announced as the vendor and the NeSI/NIWA team worked with Cray and its subcontractors to install, test and configure Kupe at the Tamaki Data Centre to provide supercomputer services until all NeSI users are moved onto Mahuika and Maui in early 2018. A number of design workshops and training sessions were held and in late December all FitzRoy users were successfully migrated from FitzRoy to Kupe.

As part of NeSI's commitment to delivering fit-for-purpose and state-of-the-art computing and data services, a private cloud OpenStack environment was implemented on Kupe in late 2017. NeSI also successfully implemented virtual labs for NIWA Kupe users, which will provide a template for services on the Maui and Mahuika platforms.

To improve user support, the NeSI team designed and implemented a new method to install and run applications on Kupe.

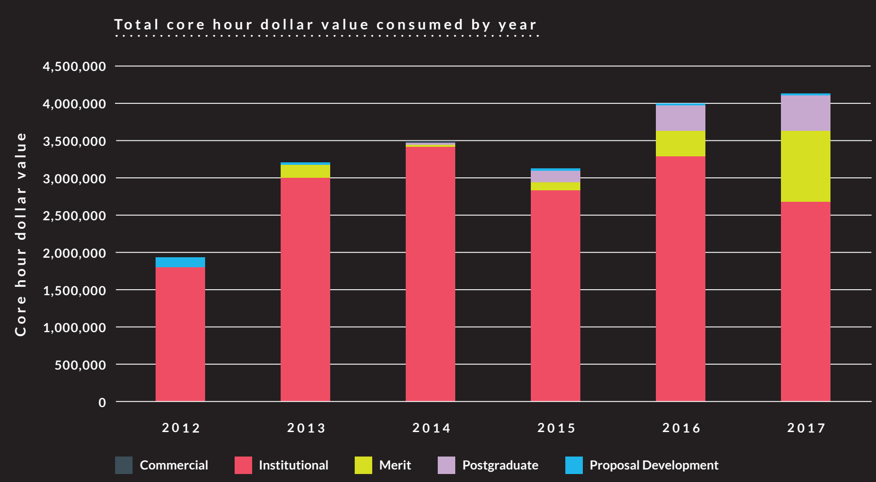

Michael Uddstrom
Platforms Manager

Gene Soudlenkov
Systems Engineer

Yuriy Halytskyy
Systems Engineer

Kumaresh Rajalingam
Analyst Programmer

Anthony Shaw
Application Support Analyst

Aaron Hicks
Systems Engineer

Jose Higino
Systems Engineer

Greg Hall
Systems Engineer

Ben Roberts
Application Support Specialist

Peter Maxwell
Application Support Specialist

Jordi Blasco
Systems Integrator

Albert Savary
Application Support

Matt Healey
Application Support Specialist

Shen Wang
Analyst Programmer

Anita Kean
Analyst Programmer

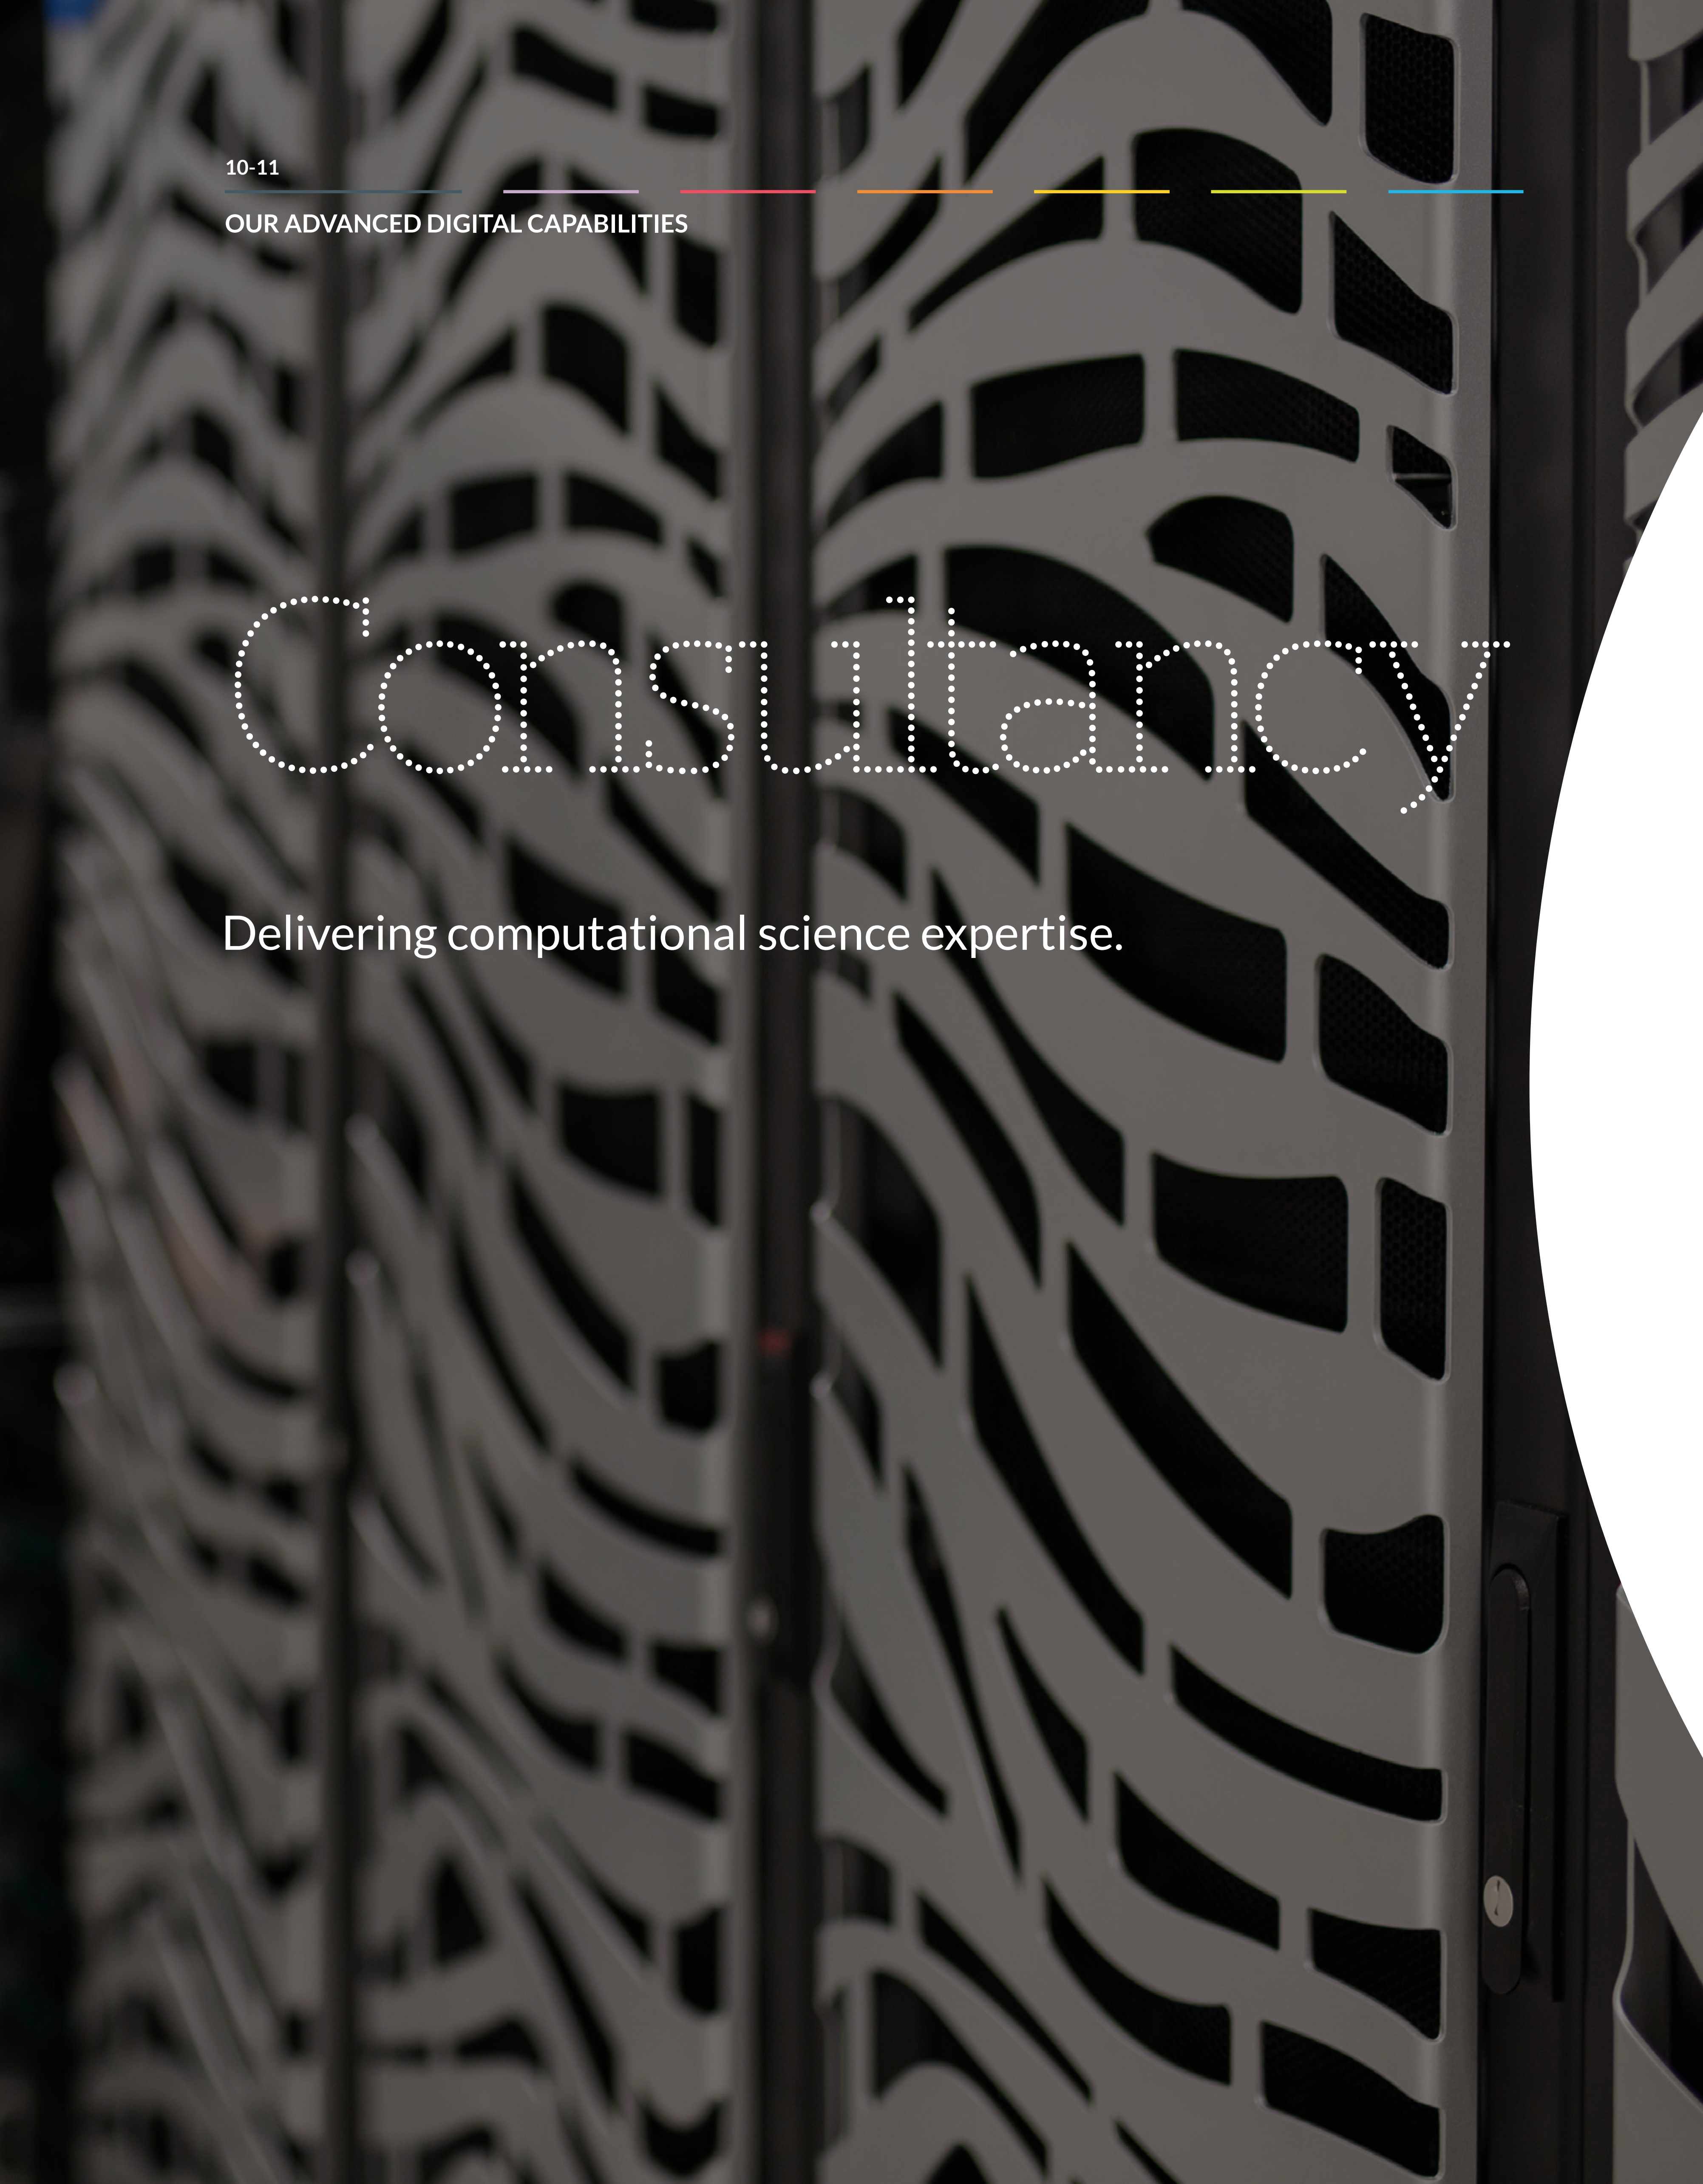

Consultancy

Delivering computational science expertise.

NeSI delivers specialist computational science expertise out into the sector, embedding NeSI's team members within research teams for a significant portion of their time. Through this collaboration with researchers, NeSI improves the performance of computational and data analytics software codes and data-driven research workflows, lifting researcher productivity and efficient use of HPC.

Lift in Consultancy capabilities during 2017:

- The NeSI Team have adopted a more proactive approach to soliciting new projects, which has successfully expanded the reach of the consultancy service across NZ. In particular, these efforts were successful in attracting new consultancy work with University of Otago, University of Auckland and Manaaki Whenua Landcare Research.
- Enhanced promotion on the NeSI website and Compute Login Interface has increased researcher awareness of the consultancy service and the level of engagement with NeSI users has grown as a result. Relationships with researchers at Victoria University of Wellington were strengthened in 2017.
- In partnership with NeSI's Collaboration and Integration Team, the Consultancy Team is working to improve the quality of the service with better tracking of projects and applications.

Chris Scott
Scientific Programmer

Wolfgang Hayek
Scientific Programmer

Alexander Pletzer
Scientific Programmer

Brian Corrie
Solutions Manager

NeSI Consultancy projects for 2017

Project name	Principal Investigator	Affiliation
Adding high performance curvilinear regridding capability to unified model pre- and post-processing tools	Hilary Oliver and Stuart Moore	NIWA
Exploring tsunami and flooding code speed-up using GPU processing	Emily Lane	NIWA
TopNet HPC optimisation	Gabriella Turek	NIWA
Using GPUs to accelerate N-body simulations of the solar system	Philip Sharp	University of Auckland
Automated objective identification of seismic phases	Ting Wang	University of Otago
Multiscale modeling of saliva secretion	James Sneyd	University of Auckland
Hidden markov chains with nonhomogeneous transition probabilities	Amina Shahzadi	University of Otago
Open CMISS on NeSI	John Rugis	University of Auckland
Scouring continuous seismic data for low-frequency earthquakes	Calum Chamberlain	Victoria University, Wellington

CASE STUDY

Shedding new light on dark matter

Richard Easther – University of Auckland

"This was a computationally challenging project and would have been impossible without the serious muscle NeSI provided. The NeSI platform tied together a collaboration between scientists based in New Zealand, Australia, Norway, and the United Kingdom."

Richard Easther, Head of the Physics Department at the University of Auckland, has been asking questions about the universe for more than 20 years. He's uncovered just as many new questions as answers in that time, but he and colleagues at Auckland and the University of Canterbury, along with a number of international collaborators, used NeSI resources to shed some potential new light on a dark part of the evolution of the universe.

"The universe is continually making space between the galaxies," says Easther. "The explosion that started our universe is still making space today."

What happened in the seconds following the Big Bang remains a source of many hypotheses and unanswered questions. In particular, researchers propose there was a period of accelerated expansion, called inflation, just after the Big Bang during which the universe expanded dramatically, before settling into a slower state of expansion.

To better understand what happened during that initial inflationary period, a research team including Easther, Grigor Aslanyan, and Layne C. Price from Auckland, and Jenni Adams from Canterbury, looked closely at the connections between the physics of inflation and the existence of ultracompact minihalos (UCMHs) – small dense clouds of dark matter that could form soon after the Big Bang.

Scientists believe that dark matter acts like supporting infrastructure or a skeleton for the universe, providing the extra gravitational force necessary to hold galaxies together and bind them to others. However, scientists have no idea what the dark matter is actually made of, and this poses a deep challenge for both particle physics and astronomy.

There are many competing ideas about what dark matter might be, and not all of them predict the existence of UCMHs. Consequently, finding UCMHs would shed light on the nature of dark matter while revealing information about the universe soon after the Big Bang, similar to the way animal and plant fossils can provide us with information about the earth's conditions in earlier geological times.

What is turning heads about this recent work is the connection it makes between the distribution of UCMHs and inflation, providing information to help develop, or maybe completely rule out, some existing inflationary models while refining our understanding of how and when we would expect minihalos to form.

Space.com published an article on this work, and it was also highlighted as an Editor's Suggestion in the Physical Review Letters (PRL) journal.

"By simulating the evolution of the perturbations during inflation, and then predicting the mass and abundances of UCMHs, they find that their non-detection is still a useful constraint on the details of inflation," says David Parkinson in PRL's summary of the findings. "This approach significantly widens the available observational pathways to understand the early Universe, and so may provide a future key piece of information as to exactly how inflation took place."

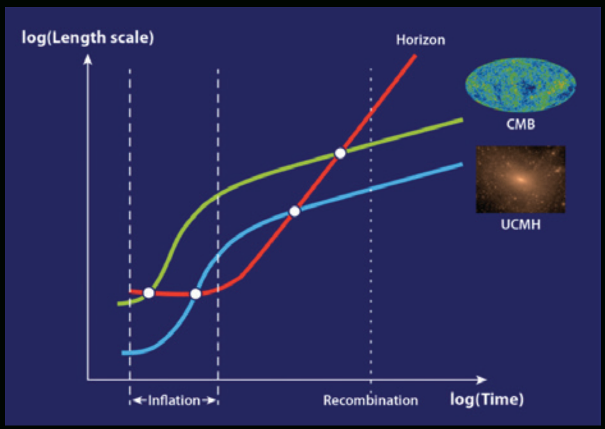

The simulations that led to these key insights were run on NeSI's high performance computing systems.

"This was a computationally challenging project, as we needed to explore the consequences of a very large range of possible inflationary models and compare those to available data," says Easther. "This would have been impossible without the serious muscle NeSI provided, and the NeSI platform tied together a collaboration between scientists based in New Zealand, Australia, Norway, and the United Kingdom."

Next, Easther's group is looking at ways in which newly released astrophysical data can further constrain models of inflation and NeSI's platforms and services will play an important role in that project as well.

"NeSI's resources makes it possible for New Zealand based scientists to do work at the forefront of modern cosmology," he says.

CASE STUDY

Cataloguing New Zealand's earthquake activities.

Calum Chamberlain – Victoria University of Wellington

Nicolas Oestreicher is pictured working at COVA, a key site used in analysing LFEs, located above the Welcome Flat hot springs in the Copland Valley in the central Southern Alps.
PHOTO CREDIT: CALUM CHAMBERLAIN

At Victoria University of Wellington, researchers are using NeSI supercomputing resources to learn more about fault zones and what their physical properties can tell us about how and when earthquakes happen.

In one project, Dr Calum Chamberlain, a Postdoctoral Research Fellow in the School of Geography, Environment and Earth Sciences at Victoria, has been building a continuous record of Low-Frequency Earthquakes (LFEs) on the deep extent of the Alpine Fault. Originally from England, Chamberlain says he chose to pursue his studies at Victoria specifically because of New Zealand's earthquake activity and the international reputations of the staff and students based in Wellington.

The Alpine Fault, stretching along the length of New Zealand's South Island, is a prime specimen for careful study as geological data suggests the fault is late in its typical ~300-year cycle between large earthquakes. Chamberlain's work building this catalogue of LFEs was part of a larger Marsden-funded research project led by seismologist Professor John Townend.

"Because we can analyse such large datasets with NeSI, we can generate large catalogues of earthquakes and look at earthquake occurrence over much longer periods of time," says Chamberlain.

The more data seismologists look at, the easier it is to identify trends or noteworthy events. In Chamberlain's case, he's looking for insights into how earthquakes impart stresses on nearby faults and how this affects deformation over time.

"The better we understand the states of stress for existing faults, and in particular, how seismic activity influences the stress state of that fault and the surrounding crust, the better we can prepare for future earthquakes," Chamberlain explains.

Using an existing network of seismometers, the Victoria team has collected a high-quality dataset from the region. Chamberlain was using a Python package he wrote for the detection and analysis of repeating and near-repeating earthquakes, however, it wasn't running as efficiently as he needed to expand the catalogue to its full potential.

"Initially, Calum relied on Python multiprocessing to conduct multiple cross-correlations in parallel, but we encountered problems due to memory usage and the correlations wouldn't run," says Chris Scott, a member of NeSI's Solutions Team. "To make efficient use of multiple processors, Calum re-wrote the internal cross-correlation routines in C, but he needed assistance with the parallel implementation of these functions. I helped him with those final steps to get their correlations to perform faster and to run with more templates at once."

Before connecting with NeSI, Chamberlain says, he and his colleagues were waiting up to four months for their computational jobs to complete on smaller local clusters. When they started using NeSI's Pan supercomputer, that wait time dropped from months to a week. Then, following Scott's consultations, they were able to run five times faster using 93% less memory on the system.

"Working with Chris, it brought it down to a matter of hours," says Chamberlain. "He helped us reduce the memory components and make the parallel architecture so much more efficient."

Chamberlain's Python package, EQcorrscan, is available on Github and a manuscript based on the work, "Repeating and near-repeating earthquake detection and analysis in Python", has been published in Seismological Research Letters. He also recently presented his findings at the 2017 Geoscience Society of New Zealand Annual Conference.

> "Because we can analyse such large datasets with NeSI, we can generate large catalogues of earthquakes and look at earthquake occurrence over much longer periods of time."

The coming year is promising to be busy as well, as Chamberlain has other projects underway with Townend and GNS Science researchers Matt Gerstenberger and Emily Warren-Smith investigating aftershocks, including in the Kaikoura area. Chamberlain has been awarded further funding by the Marsden Fund of the Royal Society of New Zealand and the Earthquake Commission (EQC) to apply his methods on a nationwide basis and develop new tools for responding to future large earthquakes.

"If I hadn't done this work with Chris, I wouldn't have been in a position to apply for those grants," notes Chamberlain. "The speedups achieved from the consultations with NeSI will greatly benefit both my upcoming Marsden and EQC projects."

Training

NeSI shares expertise and builds capability in research communities and institutions, growing digital research skills and improving researchers' abilities to make use of advanced digital capabilities. NeSI focuses on sustainability by embedding training skills within institutions and communities, taking a leadership role in growing capacity and changing research culture.

Lift in Training capabilities during 2017:

- NeSI directly contributed to building expertise and capability in NZ research communities by hosting an NZ-wide Software and Data Carpentry Instructor Training workshop in January. This event supported 27 new instructors trained from the following organisations: Plant & Food Research, University of Canterbury, Victoria University of Wellington, University of Otago, University of Auckland, MBIE, Massey University, NZGL, GNS Science, ESR, Lincoln University, and NIWA. Another instructor training session is planned for April 2018, hosted by Victoria University of Wellington.

- NeSI engaged a leader in Digital Skills training to provide an external review of its training strategy to ensure it addresses the current and future research community's needs.

- NeSI supported or participated in a number of training events across New Zealand to help researchers build skills in computational science. These events included:

  - Hands-on 'Introduction to HPC using NeSI resources' workshops,
  - A 'Data Science: When Big Data Matter' workshop
  - Multiple Software Carpentry courses co-hosted with partners such as Plant and Food Research, Manaaki Whenua Landcare Research, and Scion.

Aleksandra Pawlik
Research Communities Manager

235 RESEARCHERS TRAINED

41 EVENTS

16 LOCATIONS

KEY TOPICS DELIVERED

- Introduction to HPC Using NeSI
- Software Carpentry
- Data Carpentry

SPECIAL EVENTS

- eResearch New Zealand
- Research Bazaar 2017
- Science Coding Conference
- Workshop on Sustainable Software for Science
- RNAseq and ChipSeq workshops
- Data Science: When Big Data Matter

Training attendees

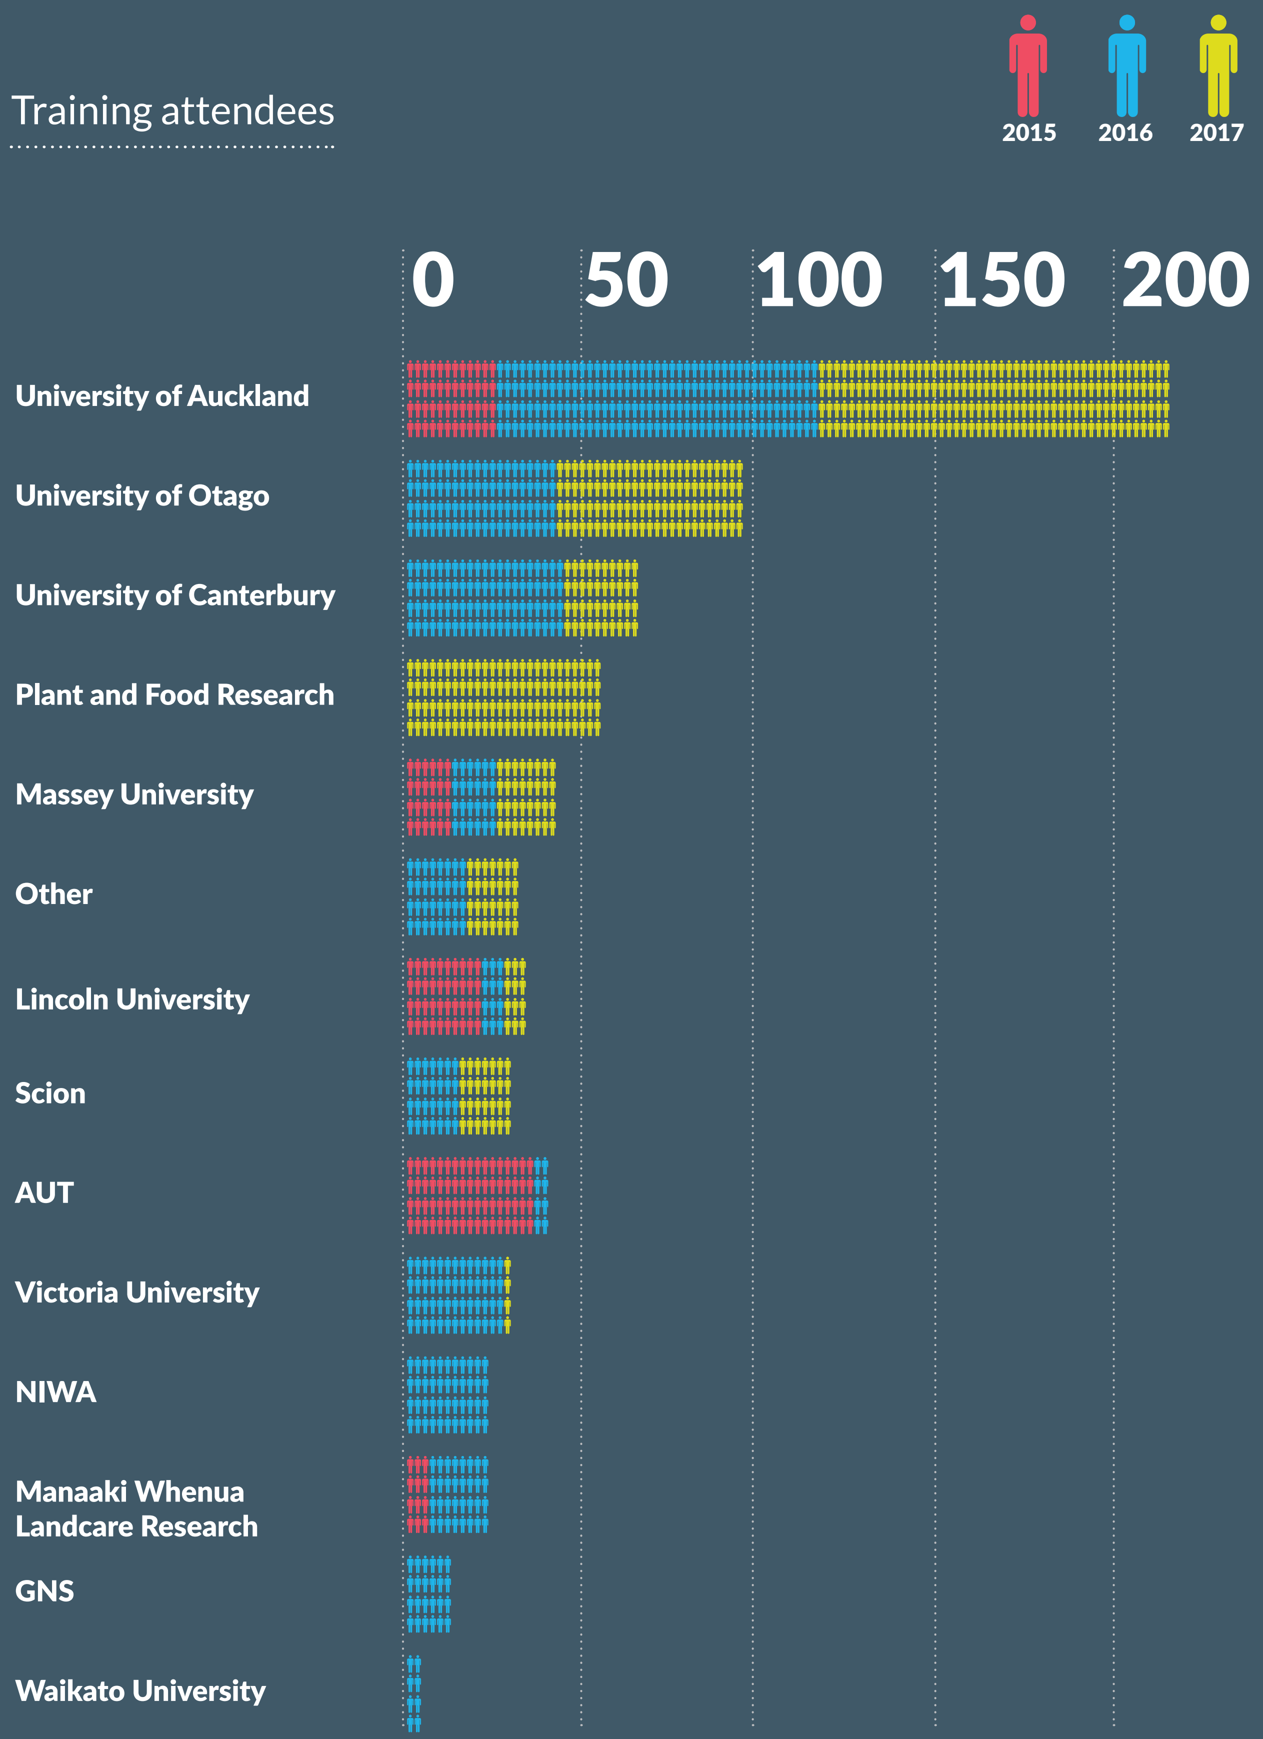

Instructor training attendees

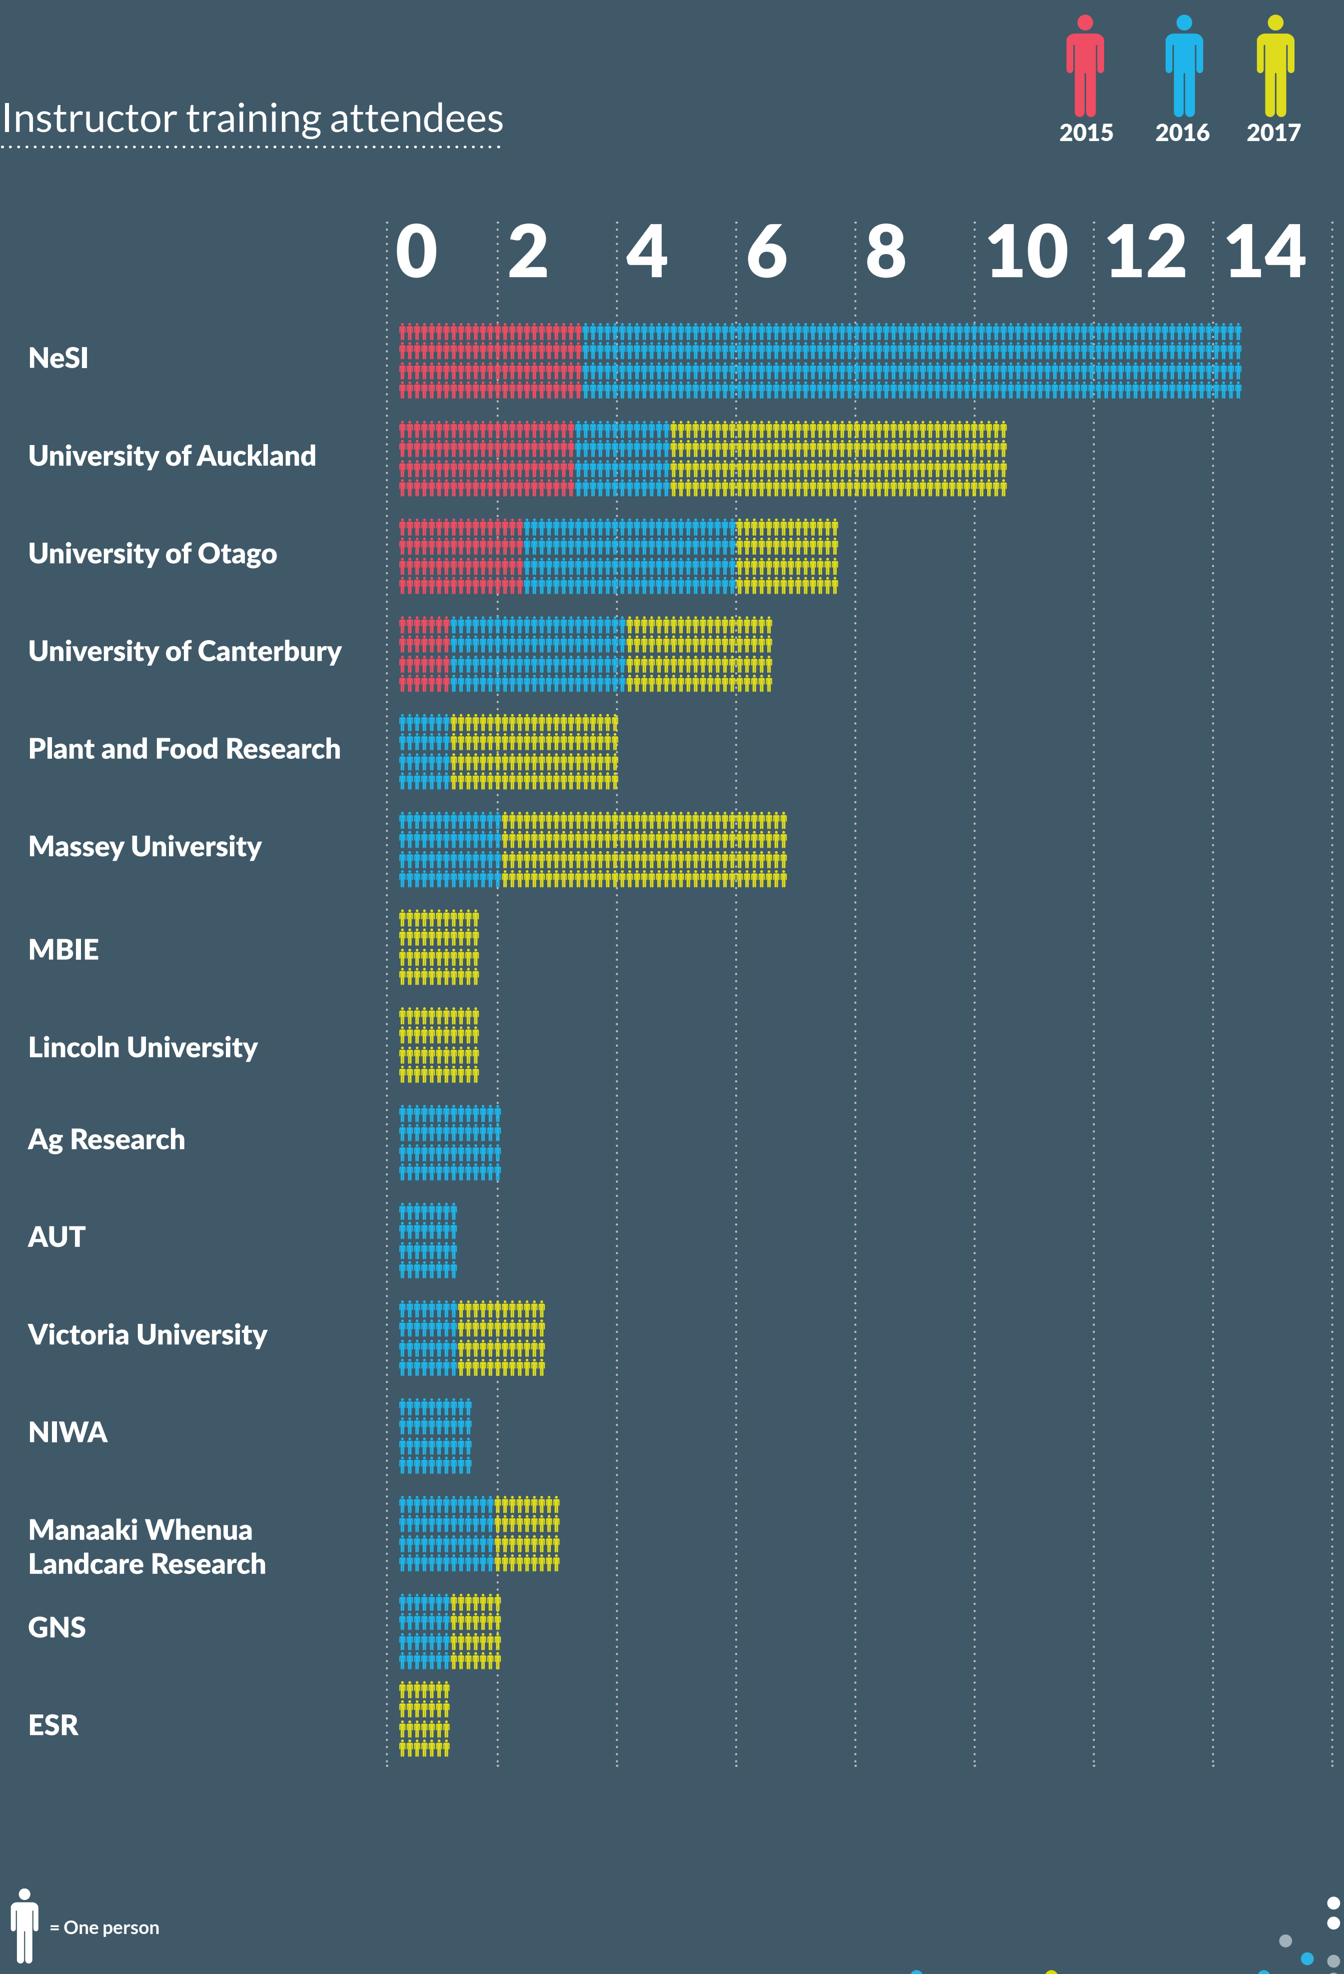

"I know I can write to NeSI whenever I need to. They're incredibly intelligent, incredibly supportive, and the turnaround time on solutions is just wonderful."

Next level HPC: NeSI user wins access to one of the world's top supercomputers.

Krista Steenbergen – Massey University

A Massey University researcher has hit the big league in the world of high performance computing (HPC), and says her experience with NeSI played a key role in helping her get there.

Krista Steenbergen is a Postdoctoral Fellow at Massey University's Centre for Theoretical Chemistry and Physics (CTCP). In April 2017, the Partnership for Advanced Computing in Europe (PRACE) awarded Steenbergen 18.1 million computing hours on one of the world's top supercomputers, the Marconi KNL system based at the CINECA supercomputing facility in Italy.

Her project collaborators include Massey colleagues Peter Schwerdtfeger and Elke Pahl, as well as Florent Calvo from the Université Grenoble Alpes in Grenoble, France. Their project was among the final 59 selected from 113 applications to PRACE's 14th Call for Proposals for Project Access, and one of just two allocations awarded to projects outside of Europe.

Marconi is a Tier-0 system, co-designed by Cineca and Lenovo, with computational power of about 20 PFLOPS and a data storage capacity of more than 20 petabytes. To put that into perspective, using an example from Indiana University, in order to match what a 1 PFLOPS system can do in just one second, you'd have to perform one calculation every second for 31,688,765 years. In terms of storage, to hold just 1 petabyte of data you would need a stack of single-sided DVDs 290 metres tall, which is just 38 metres shy of the height of Auckland's Sky Tower.

"In general, the class of HPC like Marconi is on a different level," says Steenbergen. "It's been built and configured assuming massively parallel applications. Even the way it's run assumes that."

This specialised environment is partly what inspired Steenbergen's team to pursue their project.

"My two supervisors and I had been talking for a while about wanting to do a certain kind of simulation, and we knew that both our local cluster and the NeSI resources just wouldn't be able to do that, so that's where that PRACE allocation came from. I went out and I started looking for the types of resources that allow us to do that [simulation] at that scale, and that took me overseas."

Armed with Marconi's computing power, Steenbergen's team is going to complete a set of simulations melting bulk mercury at a series of increasing pressures. They will compare the results from relativistic and non-relativistic models, and explore how relativistic effects alter or dictate mercury's material properties under pressure.

Mercury's thermodynamic properties have long been of interest to experimental and theoretical researchers as it is the only elemental metal that exists in the liquid state under standard conditions. Recent simulations have shown that mercury is a liquid at room temperature due to relativistic effects. To date, the interplay of relativistic effects, temperature and pressure have never been theoretically explored.

The work of Steenbergen's team over the next year aims to change that.

"It's going to give us a lot of information that just isn't known and really needs to be known," Steenbergen says.

For example, most studies of mercury assume a 0 Kelvin (absolute zero) temperature to reduce the computational requirements and remove a number of potentially complicating effects, such as thermal effects. By studying mercury at temperature and systematically dialling up the pressure, Steenbergen's team is expecting much more realistic results that could be ground-breaking and offer transferable insights into the properties of other important metals, like gold.

This leading edge approach was a big factor in why PRACE selected the project, but it wasn't the only reason. Feedback from the selection committee specifically mentioned Steenbergen's "excellent track record" with using HPC.

"That excellent track record in HPC experience is entirely the time that I have spent with NeSI," she says.

Steenbergen's NeSI use began in 2009 at Victoria University of Wellington, where she was pursuing her PhD. An IBM Blue Gene – one of the most powerful supercomputers available at the time – was available through NeSI and the University of Canterbury. Her PhD supervisor, Dr Nicola Gaston at the University of Auckland, planned research that incorporated high-level accurate quantum chemistry and physics codes to conduct specialised simulations. These computational demands went far beyond what could have been done on the smaller local clusters at the time. Without access to the Blue Gene through NeSI, Steenbergen's project would have been impossible.

"We would have needed to take a completely different route and we would have had to change the project entirely," she says. "I can credit NeSI with a big part of an incredibly successful PhD. We got a lot of great work done and it is something that my supervisor and I actually continue to collaborate on because it ended up to be such great research. I am incredibly grateful to NeSI."

Now at Massey University, Steenbergen has projects running on NeSI's Pan system as well as smaller local clusters. Not only have her NeSI allocations over the years been valuable, she says, but the NeSI support team itself has been just as important a resource.

"They're the ones in the trenches who I write to, who support me, and who work with me," she says. "I value them incredibly and I appreciate them. They have been nothing but supportive of all of my projects."

During the span of her academic career, she has used supercomputer systems in the United States, Germany, New Zealand and now Italy. At each facility, and NeSI in particular, the technical staff have offered key support by helping compile code, improve job efficiency, and answer questions big or small.

"I know I can write to NeSI whenever I need to. They're incredibly intelligent, incredibly supportive, and the turnaround time on solutions is just wonderful," Steenbergen says. "From the outside perspective as a researcher, they work their tails off and they give exceptional support. I don't think there are enough good things I can say about the NeSI team and how much I appreciate their support."

Her Marconi allocation began in April and, as with all new computing environments, there has been a slight learning curve to getting started. But this is where her previous HPC experience has given her an edge.

“For every different supercomputer you use, there are different nuances and tweaks,” she says. “They’re all slightly different, but in general, having the HPC access, knowing what the basic hurdles will be, and having that experience put me hugely ahead of the curve. I was pretty much up and running in a couple of days.”

Steenbergen’s allocation on Marconi is valid for one year and comes with a prescribed usage schedule of 1.5 million core years every month. Since the system is built to support only massively parallel applications, she is also required to use full nodes, so her calculations needed to be able to scale up. This is different to how NeSI operates, where nodes are often shared and the systems are configured to accommodate a range of requirements around numbers of processors and amounts of memory or storage.

“For a system like Pan, where NeSI is trying to meet the needs of all its different users, that’s an incredible challenge, so that’s one of the reasons why I’m so appreciative of NeSI’s support,” Steenbergen says.

Looking ahead at her year on Marconi, Steenbergen is already predicting a slew of conference presentations and journal papers.

“It’s going to be very busy and very exciting. I’m trying to make it a priority that I give back in the way of research outputs at a scale comparable to what they’ve given me,” she says. “I think we’re going to end up with some amazing science out of this and I’m very excited about the things we’re going to learn.”

Will there be room in her busy schedule for a trip to CINECA in Italy?

“It’s not in the plans yet, but there’s no question I will do it in the future,” Steenbergen says. “It would be wonderful to meet Marconi.”

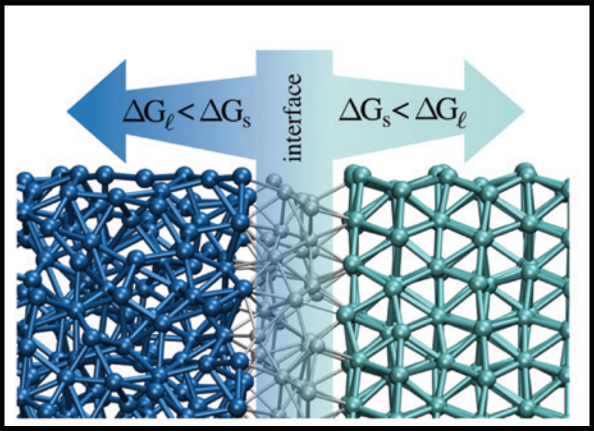

Figure taken from Krista Steenbergen’s paper, Accurate, Large-Scale Density Functional Melting of Hg: Relativistic Effects Decrease Melting Temperature by 160 K, published in the Journal of Physical Chemistry on March 16, 2017. Using first-principles calculations and the “interface pinning” method in large-scale density functional molecular dynamics simulations of bulk melting, Steenbergen and her colleagues proved that mercury is a liquid at room temperature due to relativistic effects.


CRAY
CS500
CRAY
CS500
CS500

Data transfer

NeSI supports researchers to transfer and share their research data sets, working with research communities, institutions and REANNZ to improve end-to-end performance and reduce time to solution.

During 2017, NeSI saw a continued small number of large-scale data-intensive research project proposals. The largest projects are driving rapid growth in data transfers into NeSI.

Lift in Data Transfer capabilities during 2017:

- Improvements to the service, in parallel with working with projects that are driving its heavy use, saw significant utilisation of the service at the beginning of the year. Usage of the data transfer service remains high with growing demand from Otago, LIC, and AgResearch genomics workflows.
- Improvements to the Auckland and NIWA data transfer endpoints have increased the reliability and performance of the service, enhancing the user experience for the early adopters.
- Design of NeSI's Data Transfer Solution for Kupe was completed and discussions are ongoing with REANNZ regarding how to best support and serve the network requirements. Data Transfer nodes in and out of Greta Point HPCF and Tamaki Data Centre for file replication are currently being designed to underpin the new platforms capabilities.
- NeSI completed an audit of its current Data Transfer service, infrastructure, user support, and user experience. Promoting Data Transfer to the research sector will be a priority in 2018.

High speed data transfer service – Terabytes transferred per month

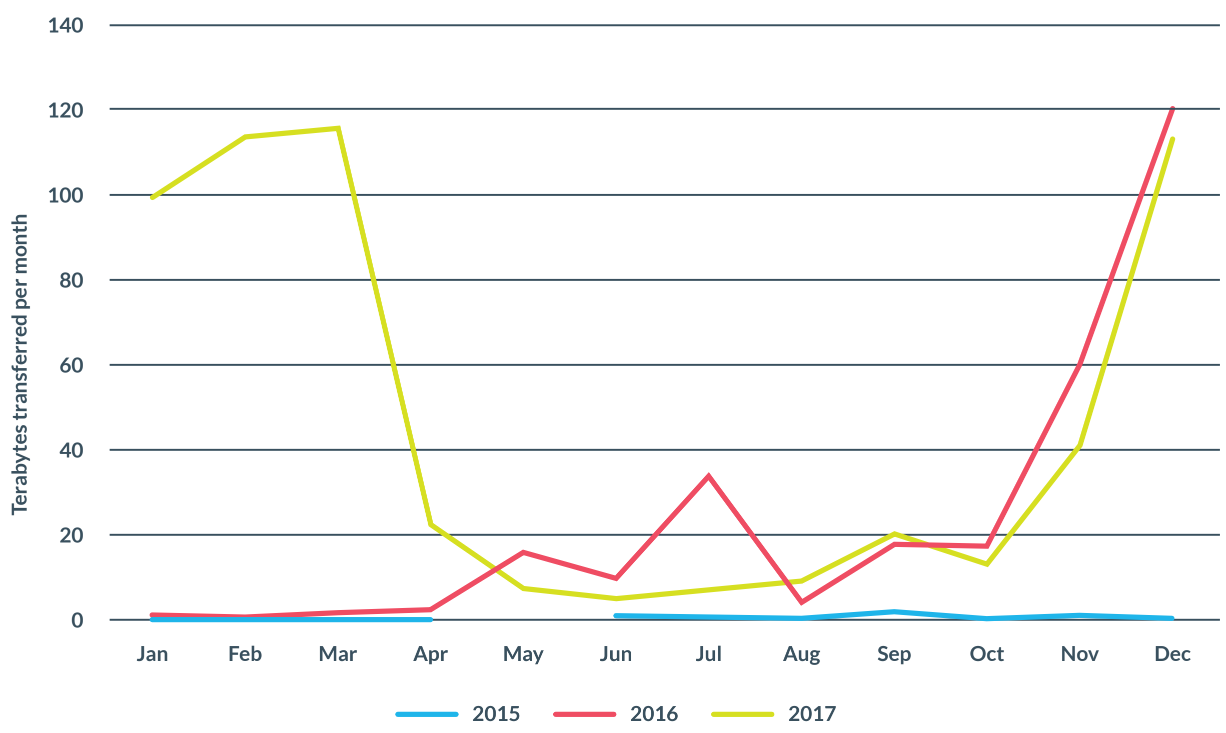

As NZ researchers take on more large-scale and data-intensive projects, we are seeing significant growth in the usage of and demand for reliable and robust data transfer solutions.

Turning data points into research insights.

James Shepherd – Manaaki Whenua – Landcare Research

"By splitting the computational workload over hundreds of cores, we were able to significantly reduce the amount of time needed to produce these data-rich high-quality models."

The more we know about an area's topography and landscape, the better plans we can make for land use proposals, transportation strategies, even natural disaster planning. At Manaaki Whenua – Landcare Research, scientists are using NeSI supercomputers to help paint a more accurate picture of the Greater Wellington Region.

LiDAR, an acronym for Light Detection and Ranging, is a surveying method that uses the light of a laser to measure the Earth's surface for building high-resolution 3-dimensional (3D) maps and models.

Using LiDAR, Dr James Shepherd's team at Landcare Research is hoping to better understand the composition of Wellington's forest canopies as well as what's found underneath the canopy, such as vegetation class or group, carbon, biomass, terrain, and impacts from possums.

Recently, the entire Greater Wellington region (8,000 sq. km), including Kapiti Island, was scanned using aerial LiDAR instruments. The scan, which required 1,266 flight lines over the region, collected approximately 82 billion points of 3D data. The dataset was massive and required powerful computing resources to extract the valuable information it contained.

This project needed a great deal of storage space. At least 30 TB, enough data to fill more than 6,000 DVDs, was required to store the LiDAR data and its intermediate files. The processing power requirements were also formidable: processing the LiDAR images on a desktop computer running non-stop would have taken more than 57,000 hours, or six-and-a-half years.

NeSI's high-performance computing facilities at the University of Auckland made it possible to store all the data and cut the total processing time down to a few days.

"Our supercomputer Pan was an ideal resource for this project," says Ben Roberts, an Application Support Specialist at NeSI. "By splitting the computational workload over hundreds of cores, we were able to significantly reduce the amount of time needed to produce these data-rich high-quality models."

The models themselves are built from point clouds, which are sets of data points that represent the external surface of objects and terrain.

"NeSI's support is critical to our research into automatic processing and classification of LiDAR point clouds, and the integration of those point clouds into digital terrain models and canopy height models," says Ben Jolly, Scientific Programmer, Landcare Research. "The data we are collecting has applications in a variety of sectors, including hydrological modelling (e.g. flood modelling), roading, air traffic approaches to the Wellington airport, and more."

Eventually, the digital terrain and canopy height models produced from the LiDAR data will be made available to the public through Land Information New Zealand (LINZ).

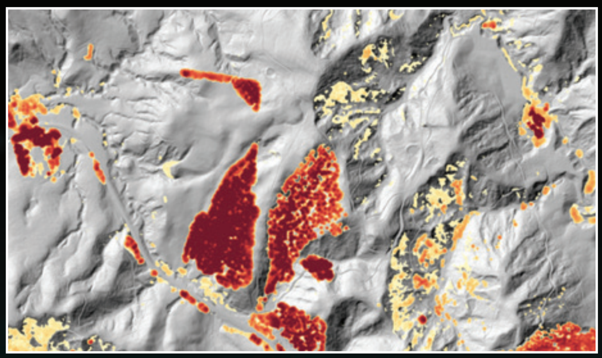

LiDAR digital elevation model with heat-map overlay indicating the height of all trees/shrubs. Photo credit: James Shepherd.

Meeting our targets.

NeSI focuses on the following objectives, tracking performance against related Key Performance Indicators (KPIs).

1 | Support New Zealand's research priorities

KPI
Annual case studies published aligned with government priorities

NUMBER	ACTUAL
≥20 CASE STUDIES	21

2 | Grow advanced skills that can apply high-tech capabilities to challenging research questions

KPI
Annually, users report significant benefit to any project from NeSI service

NUMBER	ACTUAL
>50 USERS	115

3 | Increase fit-for-purpose use of national research infrastructure*

KPI
Annual change to utilisation of national platforms

NUMBER	ACTUAL
+20% UTILISATION	7%

4 | Make fit-for-purpose investments aligned with sector need

KPI
Annually, users indicate that NeSI services meet their needs

NUMBER	ACTUAL
>80% USER SATISFACTION	97%

5 | Enhance national service delivery consistency and performance to position NeSI for growth

KPI
Annual availability of services

NUMBER	ACTUAL
>98% AVAILABILITY	97%

6 | Realise financial contributions and revenue targets to enhance NeSI's sustainability*

KPI
Contract-to-date, ratio of collaborator commitments to Crown contribution

NUMBER	CONTRACT TO DATE
>95%	79.4%

KPI 3 and 6

*NeSI's infrastructure assets are at the end of their operational life, with a capital equipment refresh programme underway, scheduled for completion Q2 2018. Both KPI 3 and 6 are performing below the standard as anticipated for assets at this phase of their lifecycle. Once NeSI's new systems are brought online, we will expect to be on target to meet these KPI objectives by the end of June 2019.

Objective 1:

Support New Zealand's research priorities

Georgina Rae

Engagement Manager

Building capabilities within national research collaborations

Through 2017, NeSI continued to focus on building capabilities working with other sector investments. This included working closely with Genomics Aotearoa to develop a shared understanding of their research computing requirements and contribute to a collaboration with the National eResearch Collaboration Tools and Resources Project (Nectar; Australia), Centre for eResearch (University of Auckland) and University of Otago to pilot the Genomics Virtual Laboratory (a cloud-based suite of genomics analysis tools for research and training) in workshops at University of Otago and University of Auckland.

Three key national collaborations NeSI worked closely with during 2017 are QuakeCoRE, the Deep South National Science Challenge, and the Dodds Walls Centre. Each of the groups carried out research projects requiring significant sustained HPC support.

Growing broader sector uptake and impacts

Since adjustments to NeSI's Access Policy in 2014 substituting pricing to researchers for pricing to institutions, the breadth of institutions accessing NeSI has steadily grown. 2017 saw a focus for NeSI on continuing to bring new organisations on board via the various access mechanisms.

This sustained growth indicates a broadening of sector uptake. Value is also put on ensuring new organisations are supported to optimise their use of NeSI by maintaining high support service levels, engaging with local researchers and carrying out outreach and training activities which local researchers can access. This included a concerted effort to develop a pipeline of possible consultancy projects with various organisations which had not previously accessed the service and garnered its associated benefits.

OBJECTIVES

QuakeCoRE platform usage in core hour value

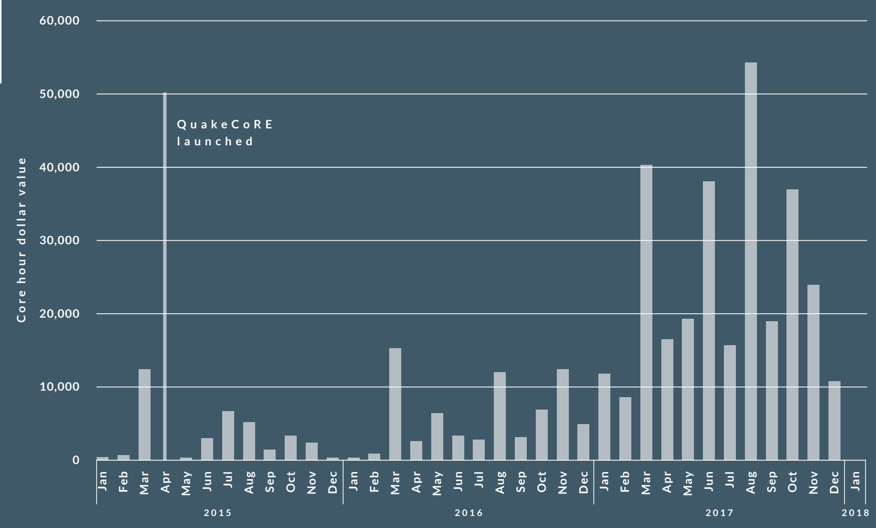

QuakeCore - NZ Centre for EarthQuake Resilience

Dodds-Walls Centre platform usage

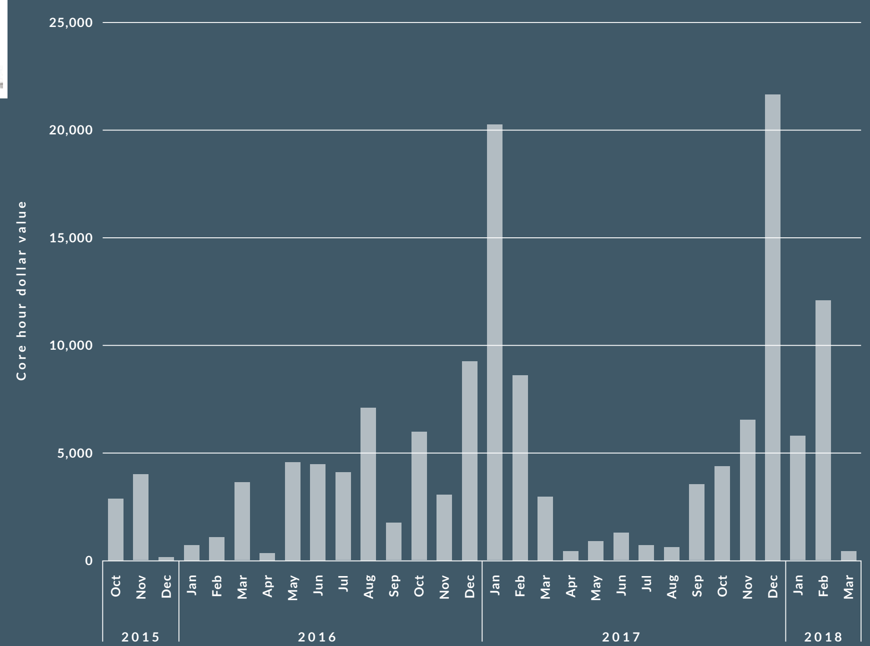

The Dodd-Walls Centre for Photonic and Quantum Technologies

Deep South National Science Challenge platform usage in core hour value

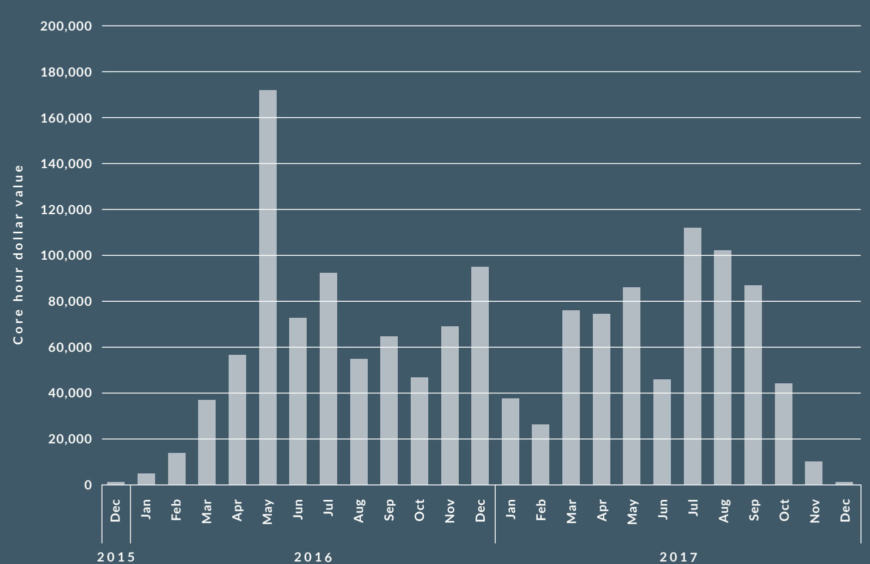

The Deep South National Science Challenge, enabling New Zealanders to adapt, manage risk, and thrive in a changing climate

Organisations that actively used NeSI platforms by quarter

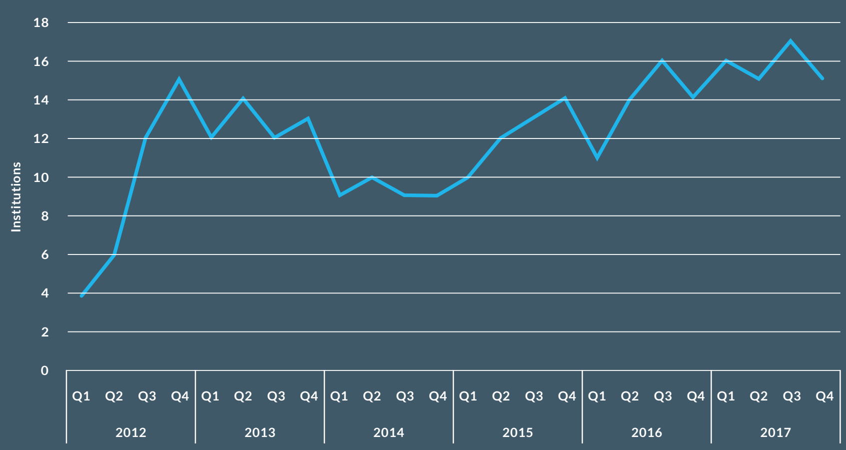

NeSI Project allocations beyond collaborator projects

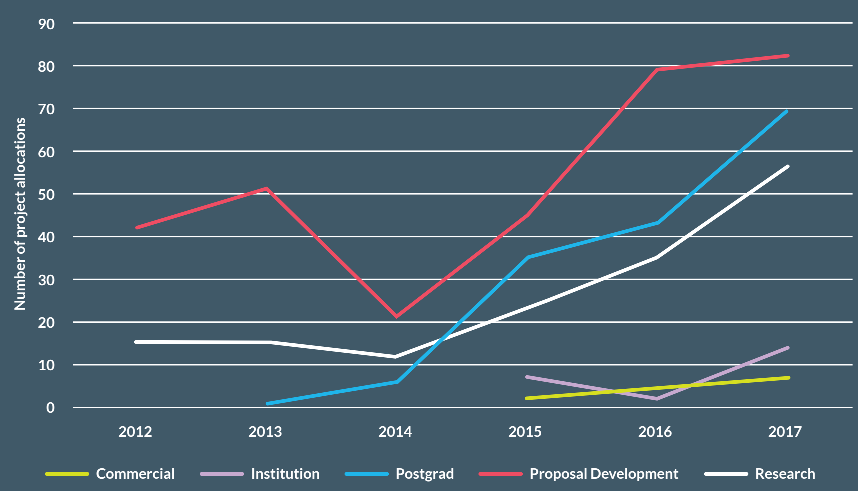

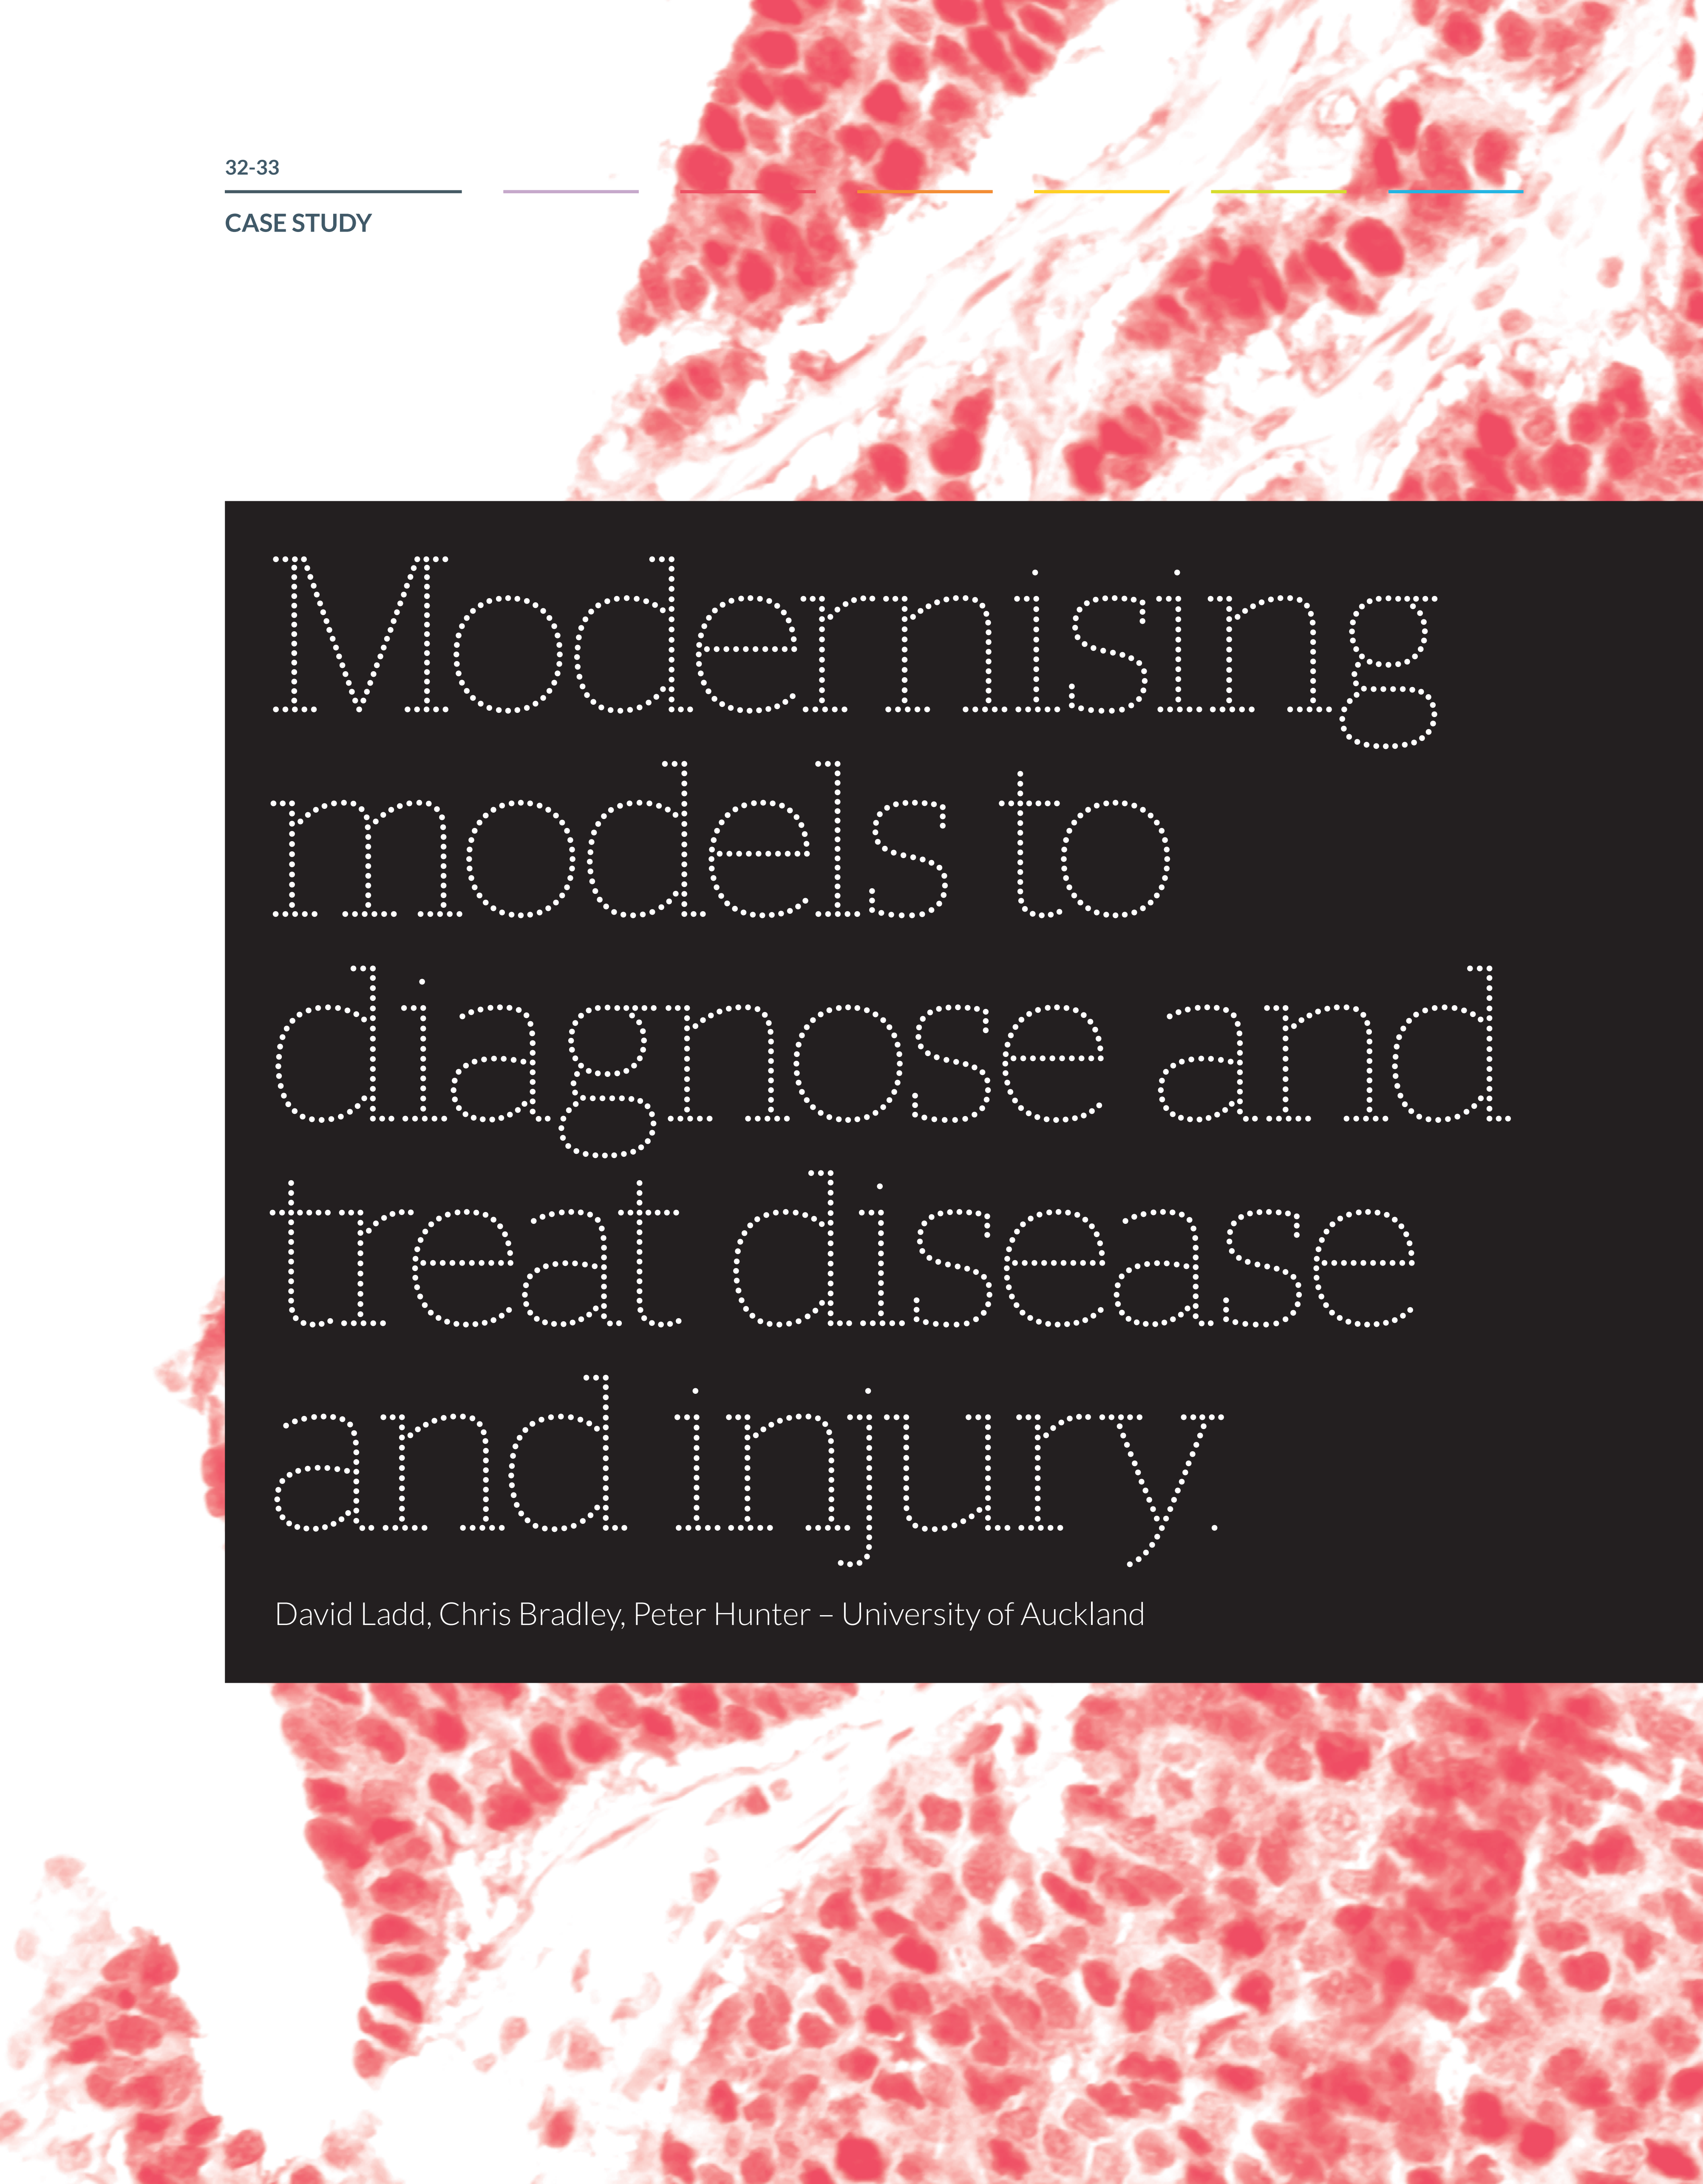

Modernising models to diagnose and treat disease and injury

David Ladd, Chris Bradley, Peter Hunter – University of Auckland

“When you’re running several big CFD problems and looking at problems on the order of hundreds of thousands or millions of degrees of freedom, you need HPC.”

Computer modelling techniques and visualisation software packages have been important resources for biomedical engineering research for decades. In New Zealand, a group of University of Auckland researchers have been at the heart of a movement to re-envision some of these tools with open-source approaches and standards-compliant practices. The results will not only enable new combinations and better linking of computational models, they will also improve researchers’ understanding of physiological functions and abilities to diagnose or treat injury or disease.

Since 2005, Chris Bradley, a Senior Research Fellow at the Auckland Bioengineering Institute (ABI), has been leading an open-source reworking of the CMISS (Continuum Mechanics, Imaging, Signal processing and System identification) programme. CMISS is a set of libraries and applications widely used in bioengineering to support computational modelling and visualisation. It was created in 1980 by Peter Hunter, the current Director of the ABI and a University of Auckland Professor of Engineering Science. Prof. Poul Nielsen and Prof. Andrew McCulloch also made major early contributions to the original CMISS packages.

Bradley’s updated version of the programme, called OpenCMISS, aims to integrate modern programming languages, data structures, and high performance hardware.

“This significant re-engineering effort represents a complete upgrade in functionality and modelling capability, particularly in terms of increased ability to optimise simulation performance on high performance, and in particular distributed, architectures,” he says.

Today, the OpenCMISS project is a collaborative effort between groups based at the University of Auckland, the University of Stuttgart, and the University of Melbourne, and is funded by both European and New Zealand research funding agencies. Developers are not only modernising OpenCMISS’s packages, they are making the tools more accessible and functional for wider use.

For example, in response to demands for better versatility across systems, Bradley and other developers have been building OpenCMISS packages to run on a variety of platforms, from Mac and Windows machines to Linux systems.

NeSI’s Pan cluster was used to support some of this work, as Bradley worked with Chris Scott from NeSI’s Solutions Team to design the underlying framework to implement OpenCMISS in a high performance computing (HPC) environment.

Another key OpenCMISS developer at the ABI is David Ladd. Ladd has been refining OpenCMISS’s fluid mechanics package, which aligned well with his doctoral thesis exploring how Computational Fluid Dynamics (CFD) models could be used for more detailed and clinically relevant analysis of fluid mechanical scenarios in cardiovascular systems.

“As computing advances have made CFD models of the vasculature more tractable, their adoption into clinical application has not been equally forthcoming,” Ladd says. “In my doctoral work, I approached vascular modelling from two sides as steps toward closing this gap.”

His project was a perfect example of a key challenge OpenCMISS hopes to address: How can researchers build models hierarchically to support better communication between the models and glean better insights from the results.

Working with patient MRI velocimetry data and using open-source and Physiomestandards compliant methods, Ladd created multiscale CFD models to improve clinicians’ ability to study a vascular system. In addition, he also created an adaptable framework for vascular fluid mechanical modelling that could be used, shared, and extended. In one application, Ladd modelled the development of a stenosis in an iliac bifurcation, where the descending aorta splits into the two main blood vessels for the legs.

“You can’t usually look at a model of an artery on its own, you also have to consider the upstream forcing mechanisms of the heart and the perfusion into downstream vascular beds,” he says. “The novel aspect of my work was that we were doing it in an open source and standards-compliant way.”

To return results in days rather than weeks, Ladd harnessed the computational power of NeSI.

“When you’re running several big CFD problems and looking at problems on the order of hundreds of thousands or millions of degrees of freedom, you need HPC,” he says. “NeSI’s resources can give researchers access to thousands of cores to compute on, versus the few that you can get from a desktop or tens that you can get from a small cluster. It’s really useful because you can run several big jobs at once.”

By the end of his thesis work, Ladd had created multiple open-source methods and solvers for the OpenCMISS library.

“We’re hoping it’s a solid base upon which we can start constructing multi-physics, multi-scale models in a standards-compliant way,” he says.

As OpenCMISS sets the stage for researchers to create more comprehensive models, other tools such as data visualization software will help researchers analyse those models on a deeper level. John Rugis, Visualization Specialist at the Centre for eResearch, knows this realm well.

“The goal of any visualization is to show 3D dynamic data in a way that is intuitively obvious,” he says. “With visualization, you know it works when someone looks at it and they don’t have to think about what they’re looking at, they’re just seeing something happening.”

Rugis recently created a dynamic 3D visualization of Ladd’s iliac stenosis simulation. A normal iliac bifurcation is featured next to one with a stenosis, making the effects of a narrowed vessel easy to see. Blood flow velocity is colour coded with blue-green indicating low velocity and yellow-red indicating high velocity. This single heart-beat animation shows the high velocity blood flow stream that results from blood vessel constriction.

“It’s showing the actual dynamic effect of that condition,” says Rugis. “It’s real science data, not just a pretty picture.”

Together, the teams and resources of the ABI, Centre for eResearch, and NeSI are helping researchers to gain broader insights and better understanding of physiological processes.

With programmes like OpenCMISS improving model creation and performance, HPC resources powering increasingly complex simulations, and data visualization techniques aiding with model analysis, researchers are moving steps closer to developing models with greater predictive capabilities and direct clinical applications.

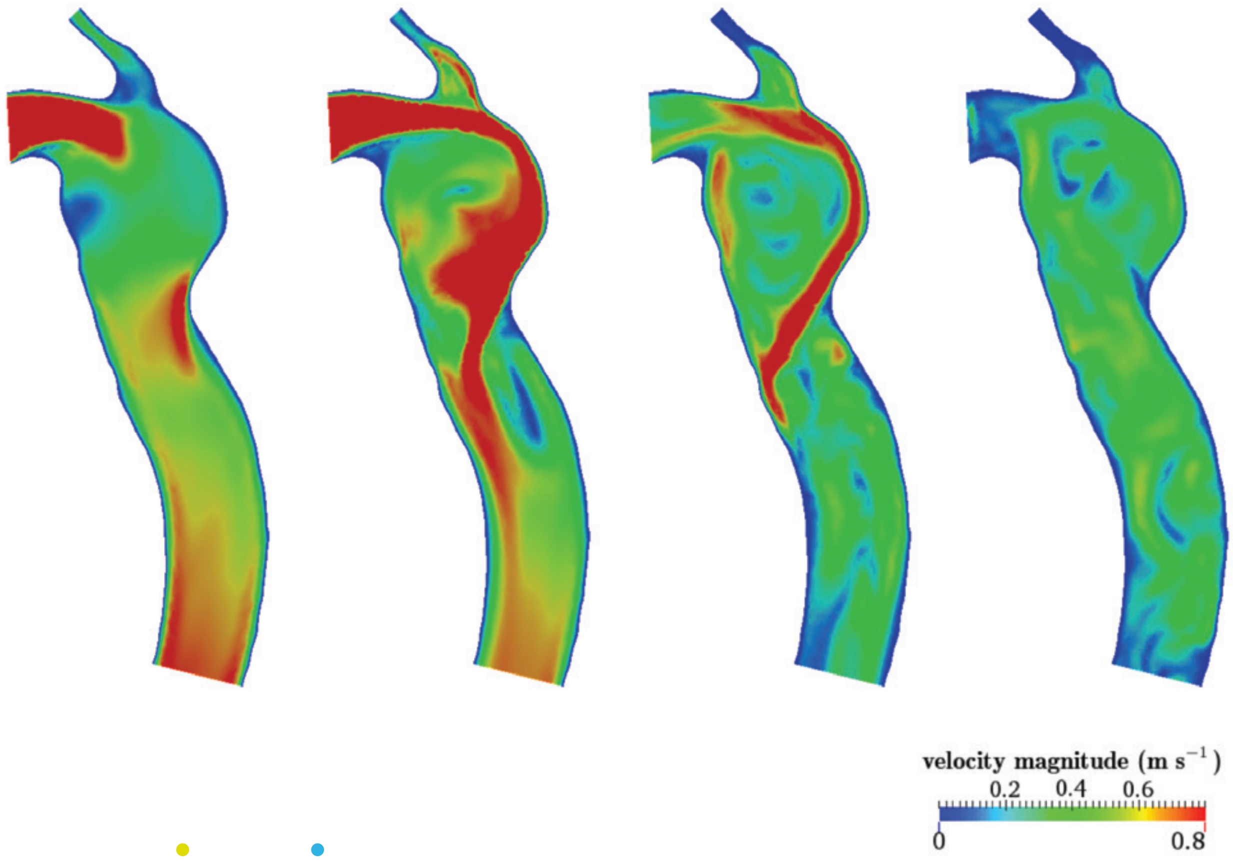

Image courtesy of David Ladd. This figure illustrates a cross-section of development of flow in an aortic aneurysm, images 100ms apart. This simulation was run on NeSI’s Pan cluster at the University of Auckland.

Embracing a diverse research sector.

NeSI continues to support a broad range of research disciplines, from the traditional big user communities of earth science and chemistry, to some of the more recent user communities such as medical science and biology.

NeSI platform usage by discipline 2012-2017

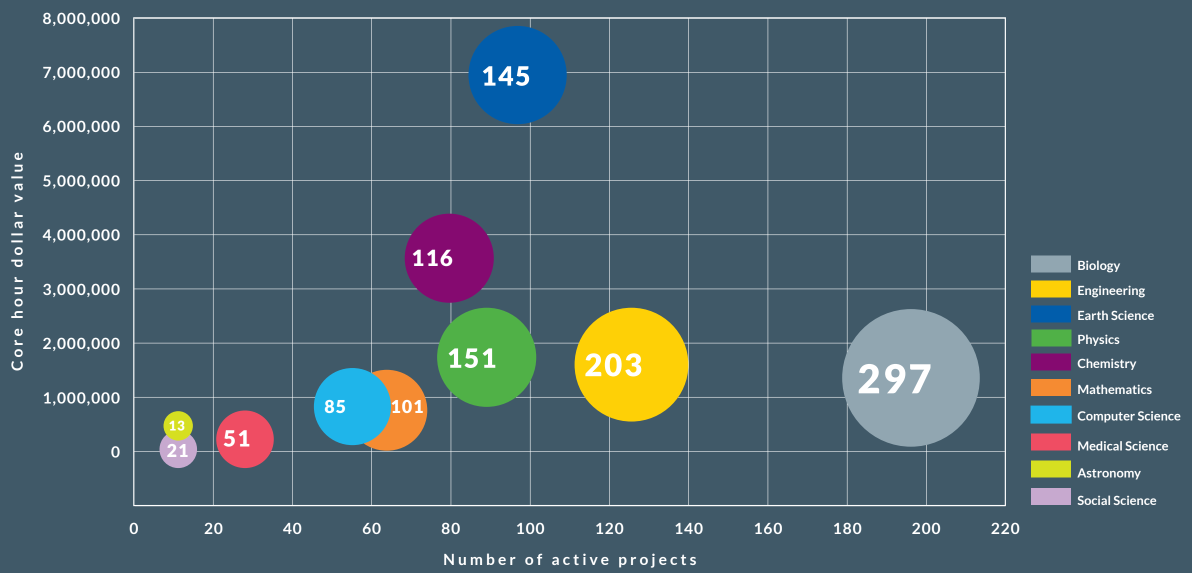

The user community is often broadly generalised into those leveraging the capability aspects of our platform (large, highly complex problems being run by users with significant experience in scientific programming) and the long-tail of researchers using NeSI's capacity functionality (lots of small, relatively simple jobs being run by users with fewer coding skills). Outside of the benefits to users of simply having access to our platforms, the capability users benefit from access to our Computational Science team via consultancy projects, while our capacity users are supported heavily by our Applications Support team via the helpdesk as well as through training events.

Research area platform utilisation

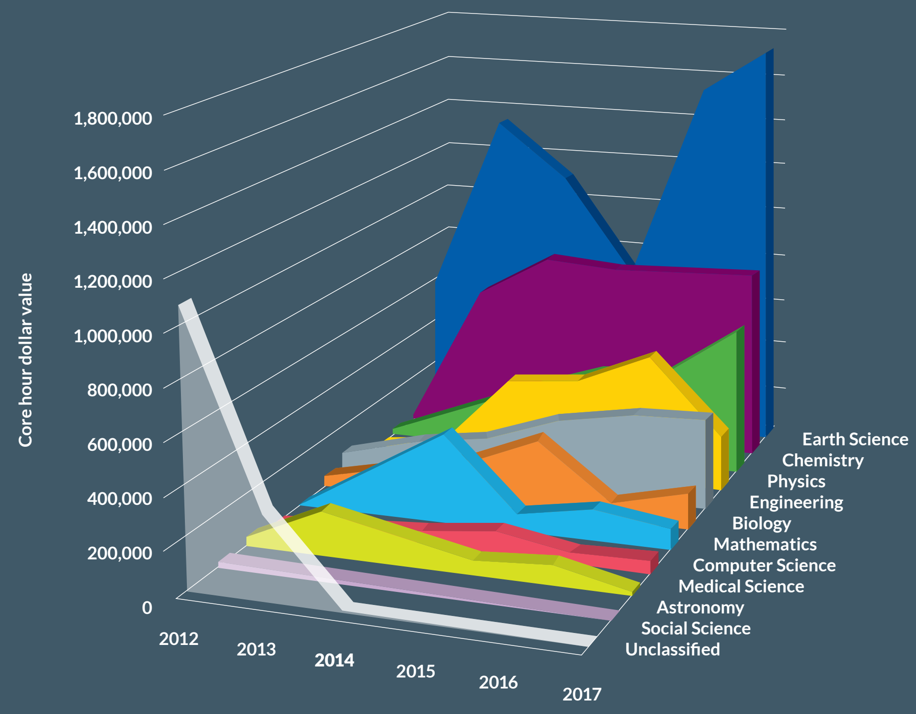

Number of new projects by discipline 2012 - 2017

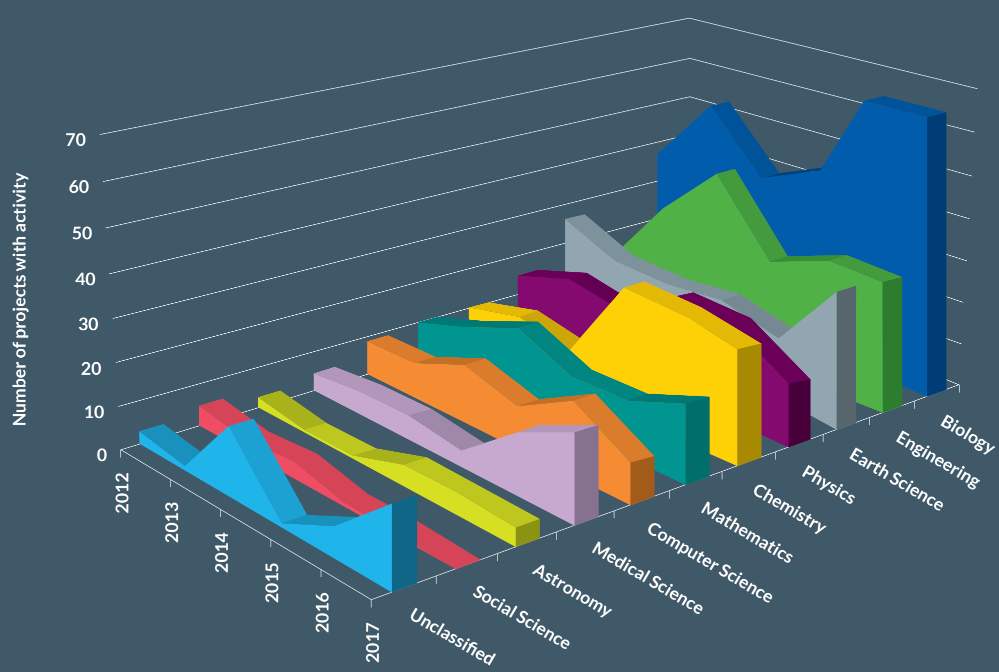

Heading in the right direction with positive outcomes for researchers.

An independent review of NeSI's achievements to date was undertaken in the first half of 2017. Reviews and evaluations are a standard part of MBIE funding agreements and this review was specified in MBIE's funding agreement with NeSI. These reviews allow MBIE and the providers to review progress, identify any strategic issues and suggest actions to improve delivery of the overall programme.

The evaluation found that NeSI has contributed to positive outcomes for researchers and research institutions, and has largely met expectations around effectiveness, relevance, impact and value for money. The financial sustainability of the programme and the business model employed could be improved, in particular, by developing strategies to attract more investment. Some options for future improvements were tabled and MBIE indicated their support for both the overall findings and the future recommendations.

KPI 1

Annual case studies published aligned with Government priorities

DEFINITION
Number of new case studies published on the NeSI website.

TARGET (Actual)	ACTUAL (YTD)
≥20	21

Twenty-one case studies were published in 2017, presenting work from all of NeSI's collaborator institutions, and including case studies highlighting research undertaken by researchers at University of Auckland, University of Otago, Victoria University of Wellington, Lincoln University, NIWA, Department of Conservation, and Genomics Aotearoa, and projects supporting the Deep South National Science Challenge and the Bio-Protection Centre of Research Excellence (CoRE).

Preparing New Zealand to adapt to climate change.

Olaf Morgenstern – NIWA and Jamie Kettleborough – UK Met Office

“The UK Met Office and NIWA have relied on NeSI staff’s deep knowledge of interpolation techniques to assess the strengths and limitations of different earth system data regridding packages.”

Imagine if you could accurately predict today what New Zealand's climate would look like in 2150. This is ultimately the aim for the Deep South National Science Challenge. A group of New Zealand researchers are contributing in an international partnership led by the UK Met Office, to the development of a new Earth System Model (ESM) that will provide policy makers and economic stakeholders with the information they need to prepare for and respond to changing climate. As extreme weather events become more frequent, flooding risks increase, and ski seasons get shorter, this and similar other models will help New Zealand communities continue to thrive in the face of climate change effects. NeSI resources are providing a portion of the computing power, and NIWA experts are helping to advise on some aspects of the design of the ESM.

Typically, ESMs contain many components (atmosphere, ocean, land, vegetation, etc.), each running on its own grid. When working with these models, it is imperative to be able to move data quickly and seamlessly across grids without introducing inaccuracies. This process of projecting data from one grid onto another is called regridding.

Regridding is becoming an increasingly important aspect of climate and weather forecasting simulations as the models move away from latitude-longitude grids, the hitherto standard for locating any position on the surface of the earth. Climate modellers have found latitude-longitude based grids inadequate near the North and South Poles, regions of high interest, because many numerical algorithms have issues when the coordinates collapse to a singular point.

Curvilinear grids such as Fig. 1 are commonly used for ESMs to alleviate the problem of singular points. However, these present their own set of challenges for regridding tools. On the one hand, it can be difficult to locate a source grid cell given a target point when grid lines are highly curved. On the other hand, earth science data tends to be attached to grid cells, and cell-centred data doesn't associate values with specific points. Instead, cell-centred data represents the average field value over cells.

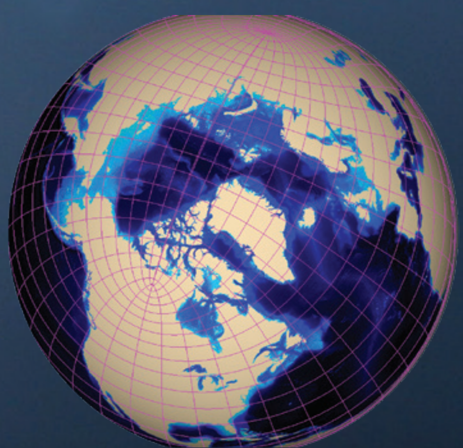

Figure 1: The tripolar grid used by the NEMO ocean model replaces the North Pole singularity with two poles located on land (Canada and Siberia).

So curvilinear grids require conservative interpolation, which involves computing the overlap of target grid cells with source grid cells. To compute the interpolation weights, one needs to compute the collisions of N source grid cells with M target cells, an N*M problem. When N and M are of the order of millions of cells, the winning regridding tools are the ones that can compute the interpolation weights in less time.

This is where NeSI's expertise has helped. Olaf Morgenstern, who is leading model development for the Deep South National Science Challenge, and Jamie Kettleborough from the UK Met Office, consulted with Alex Pletzer and Chris Scott from NeSI's Computational Science Team to evaluate different regridding tools in terms of capability, accuracy and performance.

"The UK Met Office and NIWA have relied on NeSI staff's deep knowledge of interpolation techniques to assess the strengths and limitations of different earth system data regridding packages," said Kettleborough. "The results of their study will lay the groundwork for how we prepare input data to and output data from climate model simulations, including the Deep South Challenge."

Of the tools evaluated, NeSI found the Earth System Modelling Framework (ESMF) had all the features required to regrid present and future data most effectively across a variety of curvilinear grids.

"A notable result was that ESMF conservative regridding takes only slightly more time than some existing rectilinear regridding tools, while the latter can only handle latitude-longitude grids," said Pletzer. "We also found that computing the interpolation weights was three orders of magnitudes more expensive than the interpolation proper. Every effort should be made to reduce the number of times interpolation weights are computed – we recommend the weights to be reused across different fields and time steps whenever possible."

Since other earth science researchers could benefit from these findings, NeSI plans to share them more widely. One of the first places will be Iris, an open-source Python library created by the UK Met Office to enable the visualisation of weather and climate data.

As for the Deep South Challenge, NeSI will continue to support this and other projects funded under the Challenge's Earth System Modelling and Prediction Programme. For more information on these and other Deep South Challenge news and events, visit The Deep South website.

Objective 2

Grow advanced skills that can apply high-tech capabilities to challenging research questions.

“Having the HPC access, knowing what the basic hurdles will be, and having that experience puts me hugely ahead of the curve.”

KRISTA STEENBERGEN, CENTRE FOR THEORETICAL CHEMISTRY AND PHYSICS, MASSEY UNIVERSITY

Sustained growth in users applying HPC to research

NeSI has seen a continued year-on-year increase in users on our platforms. This trend is true for users actively submitting computer jobs, as well as users accessing the system for peripheral uses such as contributing raw data and accessing the output data.

Georgina Rae

Engagement Manager

Aleksandra Pawlik

Research Communities Manager

Jana Makar

Communications Manager

Career stage of NeSI users as at early 2018

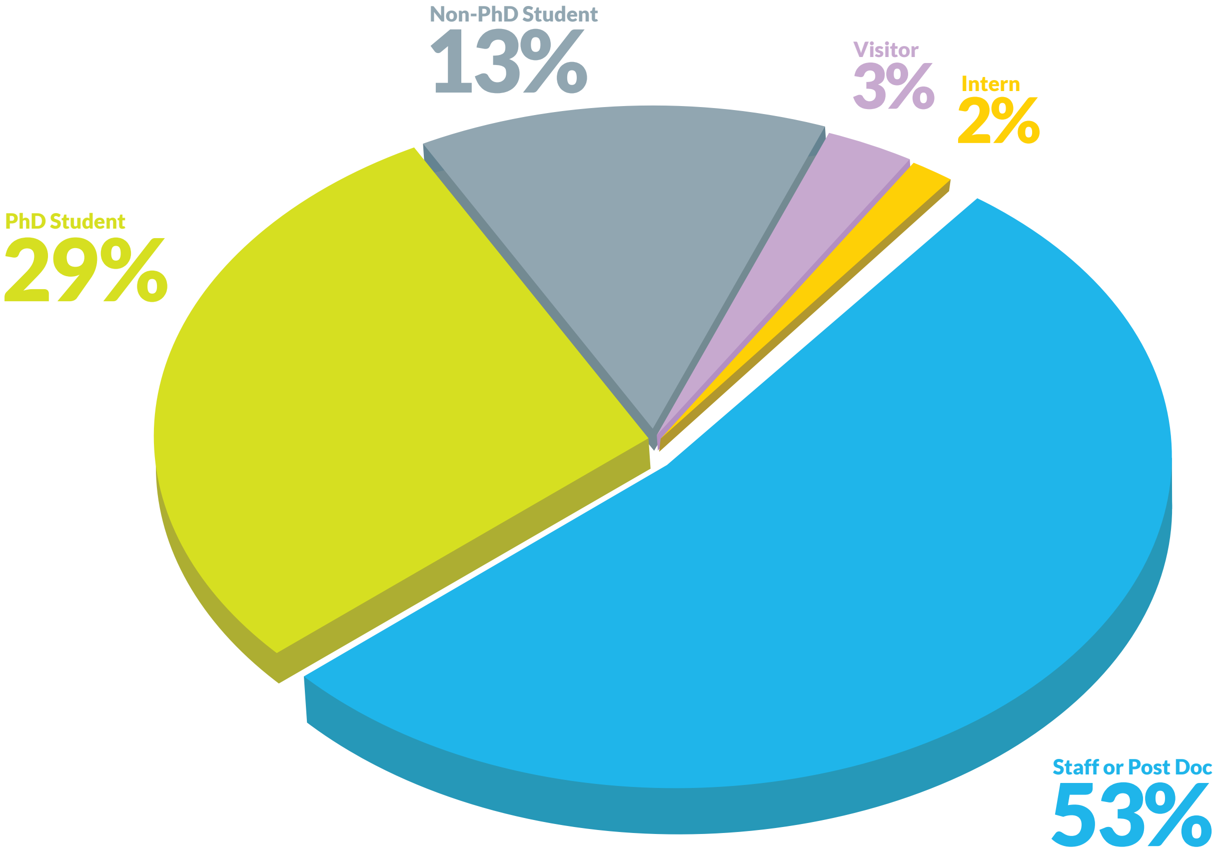

Growing researchers' digital computational literacy.

A training strategy has been in place since 2015 which aimed to help researchers develop the confidence required to use more advanced research computing resources in a scalable approach. In order to expand research computing training in New Zealand, in 2017, NeSI partnered with Software Carpentry. The partnership enabled delivery of more workshops across the country and allowed us to increase the number of instructors trained in New Zealand, strengthening and building the local training communities.

To build on the entry level digital skills training NeSI has contributed to, in 2017 an interactive 'Hands-on introduction to NeSI' workshop was developed with events held at University of Otago, University of Auckland and Manaaki Whenua Landcare Research. The purpose of these workshops was to bridge the gap between Software Carpentry and actually using NeSI, making it easier for users to start using the platform.

Overall, 15 institutions benefitted from their researchers attending a NeSI-associated training event (Plant & Food Research, University of Canterbury, Victoria University of Wellington, University of Otago, University of Auckland, Lincoln University, Massey University, AgResearch, Manaaki Whenua Landcare Research, Department of Conservation, Scion, NIWA, MBIE, NZGL and ESR), with an additional 27 instructors attending the Carpentries Instructor Training hosted by NeSI in 2017, who will go on to contribute to their own training initiatives.

NeSI's delivery against this strategy was reviewed towards the end of 2017. The findings showed that NeSI had delivered upon and exceeded the planned training efforts.

Fostering partnerships with research collaborations and institutions.

NeSI aims to not only develop and nurture its own partnerships with groups within the New Zealand and international research communities, but it also puts significant value in enabling other groups to develop such relationships. NeSI was particularly active in this space in 2017.

NeSI continued its ongoing support for the annual eResearch NZ meeting in Queenstown, co-hosting the event with NZGL and REANNZ. NeSI also contributed to the inaugural Data Science: When Big Data Matter meeting co-ordinated by Plant & Food Research and attended by a broad cross-section of Lincoln-based research organisations and CRIs nationally. Later in the year, NeSI hosted its second Science Coding Conference (formerly known as the CRI Coding Conference). This conference brings together Research Software Engineers from across CRIs (and this year also universities) to share their experiences with the aim of developing a strong sense of community amongst the group. The final major event of 2017 for NeSI was hosting the 13th IEEE International Conference on eScience. The event brought together e-researchers from across the world to promote and encourage all aspects of eScience and its associated technologies, applications, and tools.

By not only attending such events but also co-ordinating, leading and partnering on events, NeSI has established itself as a key contributor to the eResearch community.

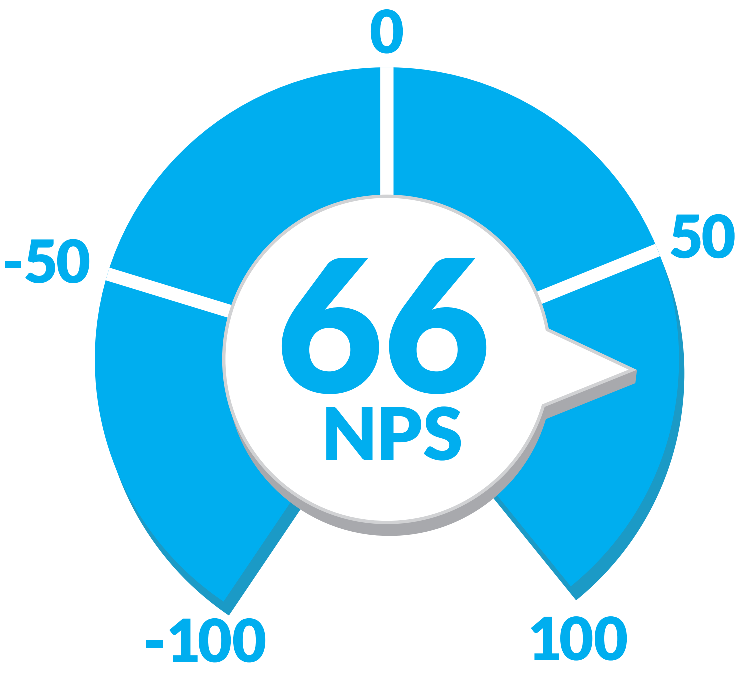

Net Promoter Score generated from responses to NeSI project survey in 2017.

KPI 2

Annually, users report significant benefit to any project from services.

115 users reported significant benefit when using the NeSI services and platforms during 2017.

DEFINITION

Number of PIs that tick the box "this project received significant benefit from NeSI" in the project closure report.

TARGET (Actual)	ACTUAL (YTD)
>50	115

NeSI's routine project survey, shared with researchers at the end of a project with NeSI, asks researchers whether they would recommend NeSI's services. Based on our survey results to date, NeSI has maintained a score of 66. This is a very high score – any positive score is seen as good, with scores between +50 to +80 seen as high performing. This shows that NeSI has loyalty and satisfaction from its user base.

> "I support the Software Carpentry initiative wholeheartedly and am glad for the role NeSI has taken with it in New Zealand."
>
> TOM KELLY, BIOINFORMATICS, UNIVERSITY OF OTAGO

Priorities for conservation of New Zealand's ecosystems and species

David Burlace – Department of Conservation

“Ben helped us start to think outside of the limitations of running this analysis on our local machine to considering the possibilities of running hundreds and thousands of iterations.”

Imagine trying to complete a jigsaw puzzle, but using pieces that constantly change their shape. Every few years, New Zealand’s Department of Conservation (DOC) has a puzzle like that to assemble in order to make decisions around conservation strategy and allocation of conservation resources.

DOC’s puzzle pieces are divided into two categories: Ecosystem Management Units (EMUs) and Species Management Units (SMUs). EMUs are places where DOC could work to maintain a wide range of ecosystems as well as the security of threatened species. SMUs are places where extra management is required for security of particular species.

It’s no small task to assemble this puzzle, especially since the characteristics of the pieces change from year to year and there are lots of different values that determine a ‘good fit’. DOC uses a specialised conservation planning software called Zonation to handle the bulk of the work. This year, DOC also enlisted NeSI computing resources and data management expertise to help tackle the challenge of processing and analysing the data.

“What the Zonation program does is, we feed it a range of information about the different species and ecosystems that occur at each management unit and we have a number of input parameters that we can tweak to provide us a ranked, hierarchical list of sites to tell us what sites are best to manage under a different range of circumstances,” says David Burlace, Technical Advisor in Planning, Monitoring & Reporting, Science & Policy at DOC.

“In order to play with all these parameters and get a feel for what settings are the optimum for our analysis purposes, we’ve been running many different scenarios with different weights on the input parameters. The computation for a single run was taking up to two hours on our local machine. Now we’re at a point where we’ve run up to 1,000 different scenarios with NeSI, which would have been completely unfeasible with DOC’s internal resources.”

The prioritisation list covers more than 500 species, and includes nearly 1,400 management units ranging in size from one hectare to more than 50,000 hectares. The candidate management units are identified by panels of in-house and external experts, and collectively they provide examples of nearly all New Zealand’s ecosystem types, ranging from coastal pohutukawa to tussock grasslands to braided rivers. Many of the management units contain several valuable ecosystem types.

Ranking is dependent on each ecosystem’s condition, which is determined by the pressures and pests known to be present, as well as the benefits of past management. Ranking also considers the potential for an ecosystem’s condition to improve if management is initiated or intensified, or lose integrity if management is discontinued or scaled back.

“Our goal is to ensure that our management intervention protects a full range of ecosystems and ensures threatened species persistence,” says Mr. Burlace. “With this sort of system, we can get an objective view to make sure that we’re accounting for everything that needs to be, and to ensure that a broad range of ecosystems and species are being managed.”

Ben Roberts, a member of NeSI’s Solutions Team, has been working with DOC’s Planning, Monitoring and Reporting team. In addition to offering technical assistance and automating parts of the analysis process, Dr. Roberts has also shared valuable advice on data management tactics.

“Ben helped us start to think outside of the limitations of running this analysis on our local machine to considering the possibilities of running hundreds and thousands of iterations. It’s really helped to broaden our critical thinking about how we’re approaching the analysis,” Mr. Burlace says. “He’s been really good at providing advice on how to tackle a project of this nature, and how to structure our workflows in order to manage the amount of data going in and coming out.”

The project is drawing to a close as the team is aiming to have a final set of rankings complete by early 2018. There are other puzzles to be solved, however, both in terms of further exploring the values represented by different Zonation scenarios and in terms of improving the quality and completeness of underlying ecosystem models. Mr. Burlace says DOC may call upon NeSI’s support to tackle some of those projects as well.

“We can write increasingly more sophisticated models but in order to do that, we need to continue to do this comprehensive analysis to make sure we’re using the correct weighted list of parameters,” he says. “Through this project, we’ve already begun to see the benefits of using the NeSI platform for some other projects we would like to get underway in the future.”

Objective 3

Increase fit-for-purpose use of national research infrastructure.

Incentivising efficient use of infrastructure investments

During 2017, NeSI has seen increased usage on the national platforms with 195 new users gaining access since 2016. Overall our usage has grown across the allocation classes with overall usage reaching the total capacity on the platforms, meaning more focus is necessary to ensure all allocation classes have their fair share of the platforms.

New infrastructure to replace the Pan and FitzRoy platforms were purchased during 2017 and installation started as the year drew to a close. This new infrastructure will deliver a significantly larger national capacity and also offer a more diverse set of capabilities to researchers in NZ. We anticipate that we will need to review the Access Policy in 2018 to consider the model for access to capabilities such as data archive where the costs and resources needed are not directly linked to the computation that are done on NeSI platforms today.

We have seen growth in researchers requesting large Merit allocations and we are expecting to see the implementation of caps of one million core hours for all Merit projects in 2018 to ensure there is a fair availability of Merit access to New Zealand researchers.

Michael Uddstrom

Platforms Manager

Distribution of new projects allocations per year by class

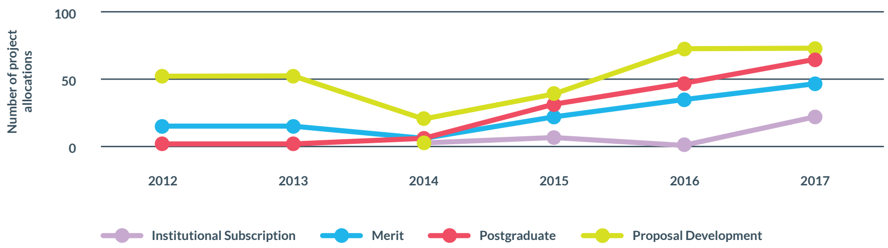

"By being able to use the parallel processing capabilities of NeSI, I was able to do the analysis within a couple of days."

TYLER BRUMMER, DEPARTMENT OF ENVIRONMENTAL MANAGEMENT, LINCOLN UNIVERSITY

Growing partnerships to increase NeSI's impact.
NeSI has received ongoing interest from organisations looking to get access to NeSI with 10 subscriptions in 2017; covering 3 with research institutions, 2 with government and 5 commercial with many new conversations that may lead to ongoing engagement into 2018. There were a number of renewals and some increases in subscriptions. Subscribers have shown much interest in what shifting across to the new platforms might offer their researchers. The new capabilities are also generating much interest, especially as NeSI becomes more data-centric, with opportunities to exploit data analytics capabilities and to build shared environments in the form of Virtual Laboratories.

Informing institutions of their researcher needs.
NeSI has been working throughout 2017 with organisations on their usage and needs. There is a growing appetite for information to support better allocation and service governance decisions. This is a key driver for prioritising development of self-service solutions for institutions in 2018.

KPI 3

Annual change to utilisation of national platforms

As the NeSI platforms have become increasingly close to being at full capacity, the growth possible has been restricted and so while we see growth of 7% for 2017 down from 20% in 2016, we expect that there is pent up demand (increasing number of users and awareness of NeSI) that when the new platforms become available mid 2018 we anticipate usage will show strong growth.

DEFINITION

Change in core-hours consumed over the 12 months to date, compared with the same period 12 months earlier (expressed as a percentage).

TARGET (Annual)	ACTUAL (Last 12 months)
+20%	7%

"Being able to perform multiple analysis in parallel on remote servers has been a huge help for this project."

STEEN KNUDSEN, SCHOOL OF BIOLOGICAL SCIENCES, UNIVERSITY OF AUCKLAND

2017 saw high utilisation and growth across access classes. With the procurement of new platforms in 2017, NeSI will offer a significant increase in capacity to the sector. Implementation started in 2017, with new platform capacities online in late Q2 2018. This will be the single biggest increase in HPC capacity that the NZ research sector has seen to date, and we anticipate as we saw in 2012 that significant growth in usage will follow this increase in capacity.

Platform usage and forecasted capacity

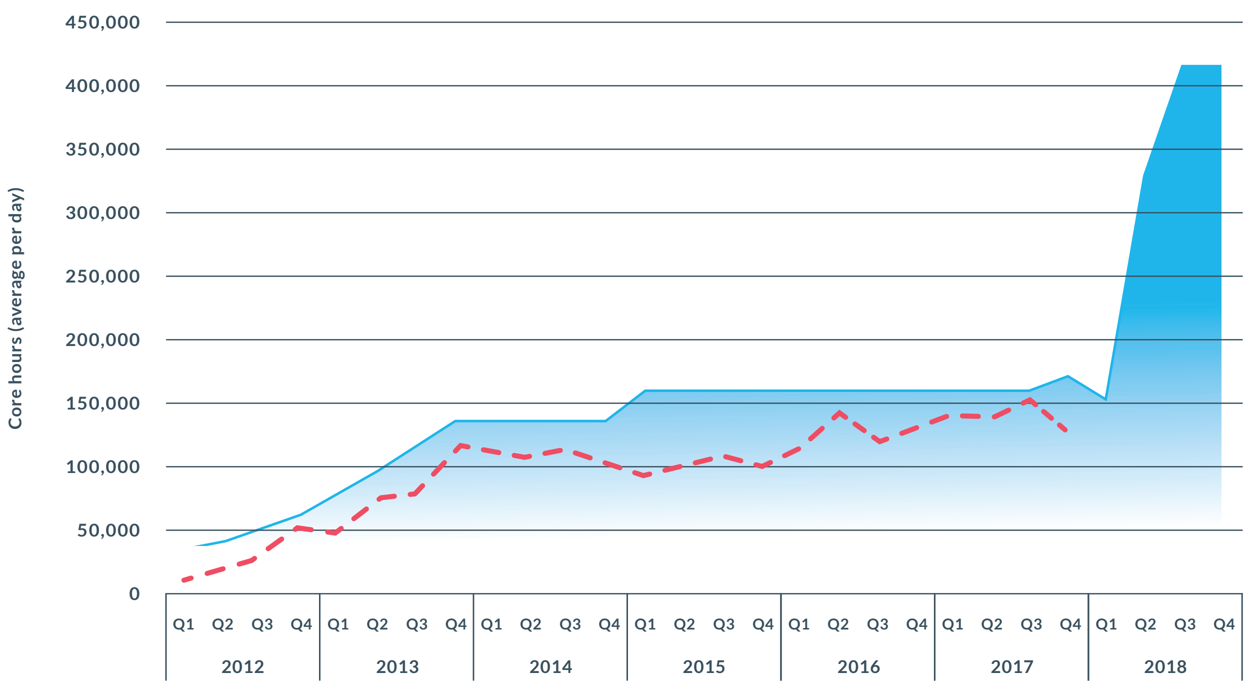

Annual utilisation by platform

FitzRoy saw record usage in its final year with the Deep South Challenge, QuakeCoRE, and Manaaki Whenua Landcare Research's Ross Sea study driving 28% growth. Pan, on the other hand, is nearing full capacity and so growth was limited. With the commissioning of Pan's replacement, Mahuika, we anticipate increasing growth across the second half of 2018.

2017 saw the final full year of usage for the FitzRoy platform with decommission set for the second week of January 2018. In December, the NIWA disaster recovery platform Kupe was made available to NeSI users as a temporary facility while FitzRoy was decommissioned and work was done on the data centre to prepare for the new Maui HPC platform. We expect Maui will bring a massive opportunity for growth with the total compute power in Maui being around 9 times that of FitzRoy.

Platform	12 months ending on 31.12.16 (CPU core hours used)	12 months ending on 31.12.17 (CPU core hours used)	% Change
Pan	36,447,914	36,852,577	+1%
FitzRoy	10,863,049	13,928,796	+28%
Kupe	n/a	19,689	n/a
Aggregated	40,197,854	51,859,280	+7%

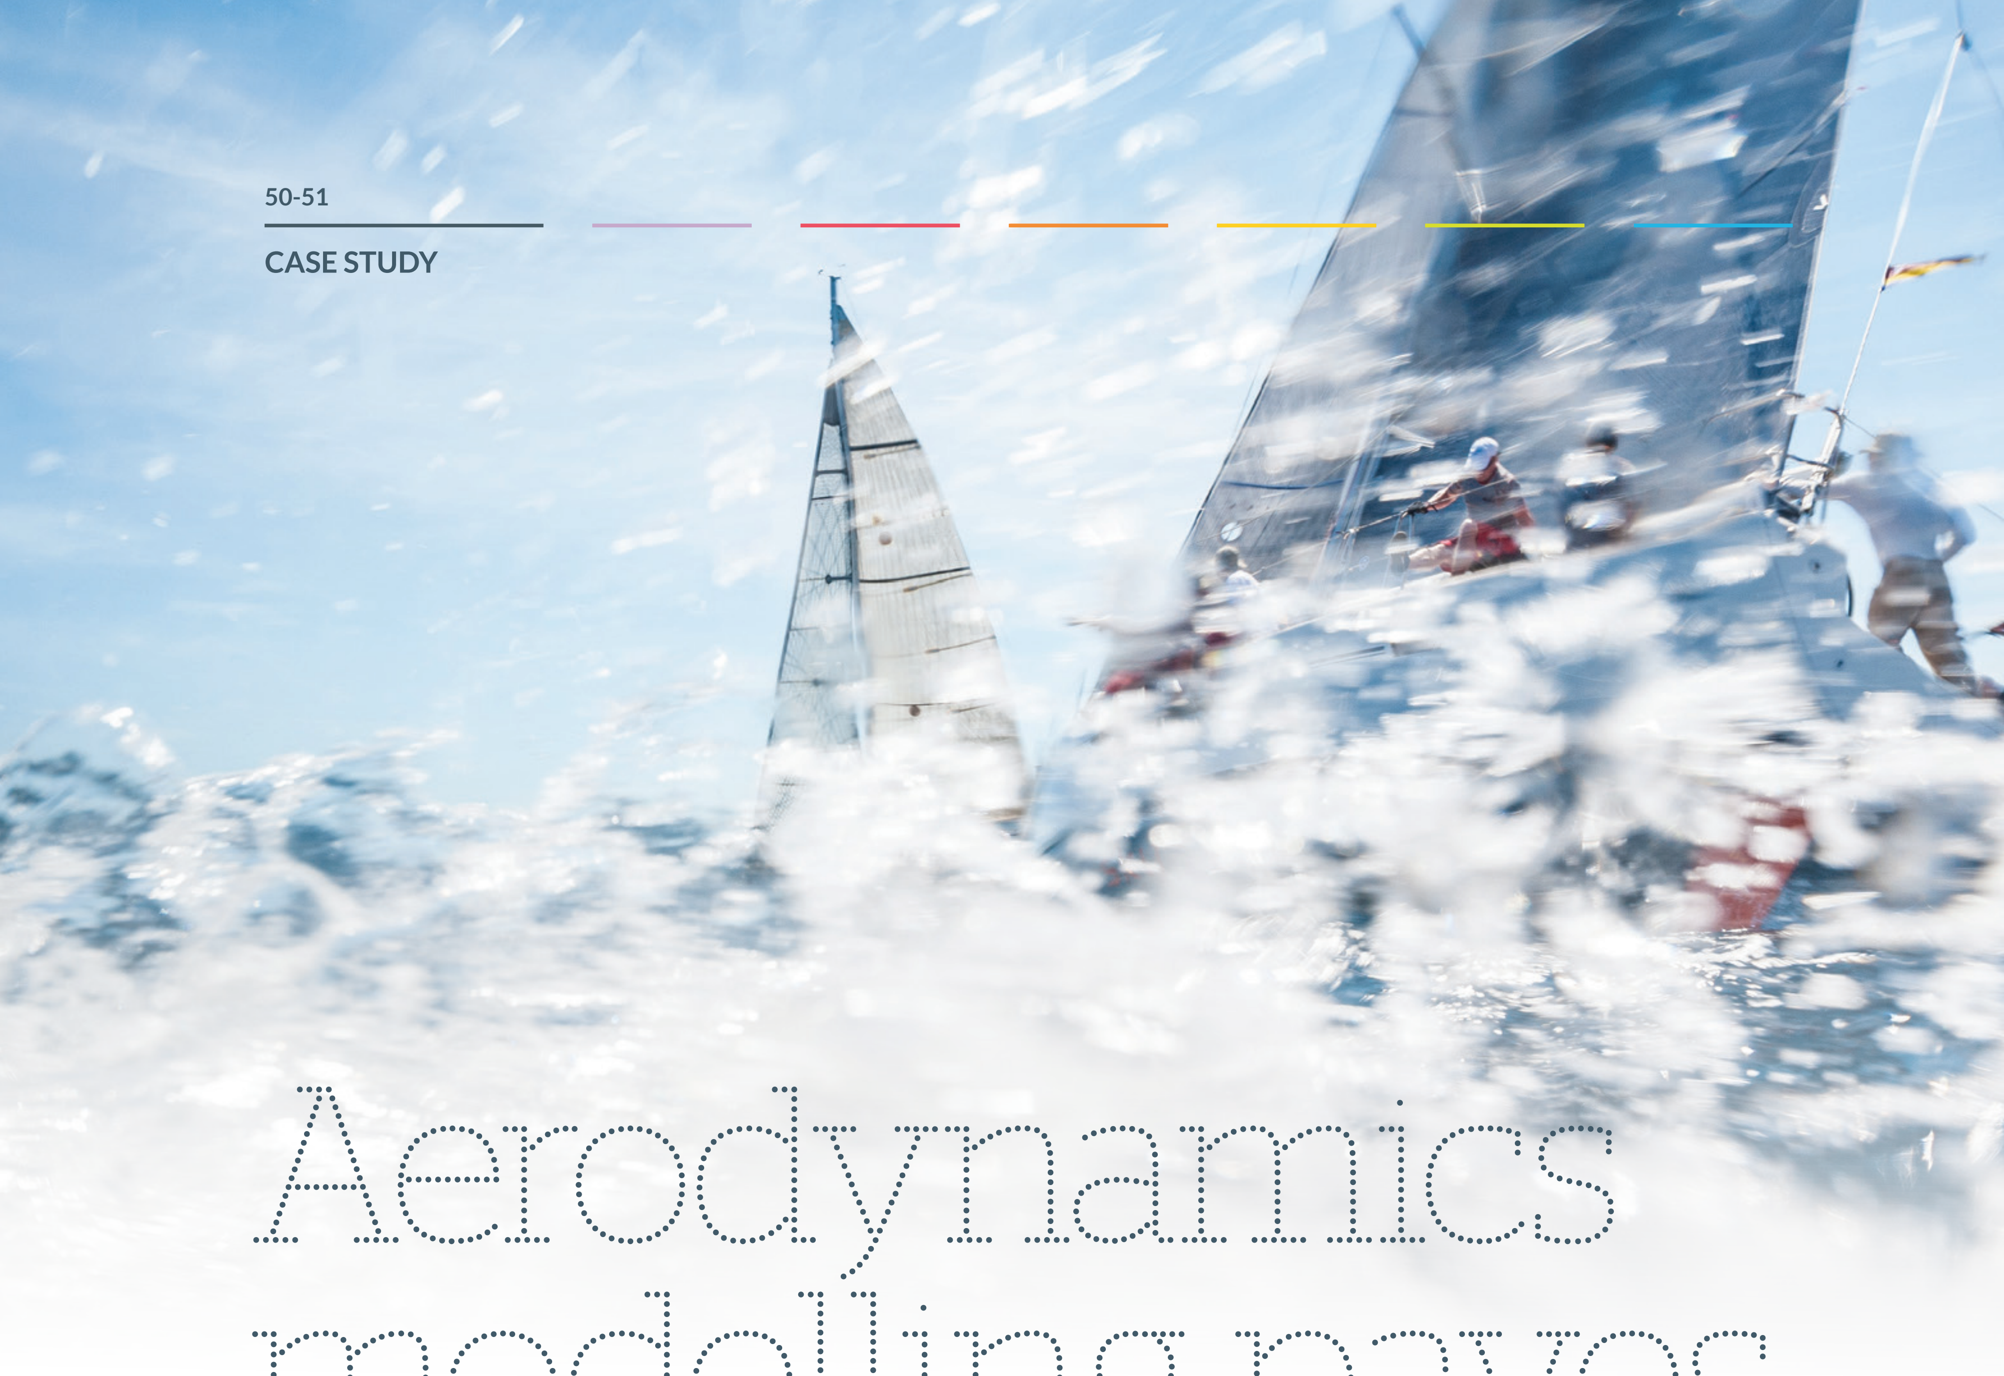

Aerodynamics modelling paves the way for improved yacht designs

Stefano Nava, Stuart Norris, John Cater – University of Auckland

“Without computational resources such as the Pan cluster at NeSI, we would be unable to use accurate, yet computationally demanding, predictive methods such as Large Eddy Simulation.”

As any experienced seafarer knows, good aerodynamics are the key to smooth sailing. In order to improve the design of sails, hulls, and masts, we need to better understand the properties of sail aerodynamics in a range of conditions.

At the University of Auckland, Stefano Nava, Stuart Norris and John Cater in the Faculty of Engineering are using NeSI computing resources to study the aerodynamics of high performance yachts.

“The fluid dynamics that characterises sailing yachts is extremely complex, due to the fact that the yacht is partly immersed in water and partly in air, with the flow being three-dimensional and turbulent,” the group notes in their latest paper, published in the International Journal of Heat and Fluid Flow. “Experimental and numerical studies have created a large body of knowledge of the physics of the problem, but at the same time have highlighted the limits to the predictive methodologies available for designers.”

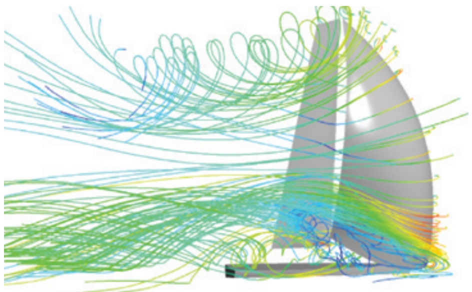

Fig. 1: Velocity streamlines of the time averaged flow around the yacht model.

Their paper compared two methods for modelling turbulent flow – Large Eddy Simulation (LES) and Reynolds Averaged Navier-Stokes (RANS) – to see if the extra computational resources required for the first method would enable a more accurate prediction of the aerodynamics of upwind sailing.

At first, the computational demands involved with this project posed a challenge. Complications with both the hardware and software components for their simulations led to performance issues and long wait times for results.

Jordi Blasco, a member of NeSI’s Solutions Team, stepped in to help. First, he created a short-term reservation on NeSI’s Pan cluster to provide a dedicated set of resources. He then found ways to improve the performance of simulation software Nava was using and the resilience of his workflow.

As a result, Nava’s computations were nearly five times more efficient and used less resources than expected, which freed up computational time on Pan for others.

“Given that Stefano had already spent weeks of computing hours on this project, the 500% efficiency improvement meant a lot, not only for him but for other users of the Pan cluster as well,” said Blasco.

Thanks to these adjustments, Nava was able to compute and collect the results he needed. The simulations were compared with experimental data previously gathered in the wind tunnels of the University of Auckland’s Aerodynamics Laboratory.

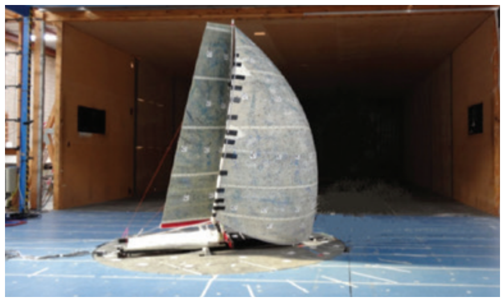

University of Auckland wind tunnel set up for a downwind sailing experiment.

Ultimately, when they compared the LES and RANS results with their experimental data, they found the LES model’s capabilities to be superior.

“In both cases, the LES model has been shown to be able to predict the correct pressure distribution for this separated flow, while the RANS simulations shows less accurate results,” their paper stated.

These results add to the growing body of work studying yacht aerodynamics. The researchers are continuing their study into the more demanding case of modelling downwind sails such as spinnakers, to be presented at the INNOV’SAIL conference this June in Lorient, France. Organised by the Cité de la Voile Eric Tabarly and the French Naval Academy Research Institute, INNOV’SAIL is an international forum for the presentation and discussion of the latest scientific and technologic research and its application in the field of high performance yachts and competitive sailing.

“Access to supercomputing resources is critical to advancing this field of study,” says Norris. “Without computational resources such as the Pan cluster at NeSI, we would be unable to use accurate, yet computationally demanding, predictive methods such as Large Eddy Simulation. It is currently not possible to use the methods on a desktop computer, due to their large CPU and memory requirements. The use of the methods enables a better understanding of the flow around yacht sails, enabling the design of faster yacht sails.”

Objective 4

Make fit-for-purpose investments aligned with sector needs.

New computing platform to power New Zealand research

Nick Jones
Director

In June 2017, NeSI announced a new integrated supercomputing platform to power computational and data intensive research in New Zealand. Science and Innovation Minister Paul Goldsmith announced the joint investment by NeSI and NIWA at the signing of a six-year contract worth $31.7 million with global supercomputing leader Cray Inc. This investment replaces and expands NeSI's national compute and data platform, extending its scope and performance to meet changing needs, and will be commissioned in early 2018.

The newly renovated High Performance Computing Facility in Wellington houses the new shared storage (left), Cray CS400 cluster (middle) and the Cray XC50 cluster (right).

The 33,500 cores in the new and enhanced supercomputing platforms will deliver up to 12 times the computing capability and more than four times the storage capacity of their predecessors, integrating three supercomputers around data intensive storage and extending this with virtualised laboratories and visualisation pipelines. The primary infrastructure will be located at NIWA in Wellington with a new replicated data infrastructure at the University of Auckland's Tamaki Data Centre providing data resilience.

Collaborative infrastructure procurement achieves significant gains.

The infrastructure RFP was completed in early 2017, leading to the announcement of Cray as the lead partner delivering the solution. As a result of negotiations, Cray Incorporated formed Cray New Zealand Ltd, leading to two international experts in high performance computing being recruited and migrating to New Zealand. Executives at both Cray and partner IBM expressed their interest in the sophistication of NeSI / NIWA's integrated design, noting it was a solution that led the way internationally in embracing the convergence of traditional HPC, big data analytics, AI, and more interactive ways of working.

Joint design delivers optimal fit for New Zealand researchers.

Following completion of the RFP the joint teams of NeSI, NIWA, Cray and IBM carried out workshops to refine the design of the new HPC systems, and the services they will provide. The features include:

- Single point of access to both the XC50 supercomputer and CS400 cluster HPC.
- Faster processors so current work is done faster.
- More processors so that more work gets done.
- GPGPU nodes to support science codes and visualisation.
- A huge memory node to support memory hungry applications.
- Interconnect performance that on the XC50 will allow jobs to scale to 1000s of processors, and on the CS400 to run very large numbers of small jobs.
- A user environment that will make it easier to manage work, develop and run research workloads/jobs, and apply data analytics tools.
- Increased storage capacity and hierarchical storage management to minimise the need to move data between the HPC storage and a user's home institution, and underpin the new interactive data analysis services.
- Vastly increased file system performance reducing the time spent reading and writing data to the file systems.

Beyond supporting NeSI's core HPC service, the new systems will deliver additional features such as:

- Interactive analyses and exploratory visualisation supported by high performance data.
- Pre- and post-processing using specialised large memory nodes, GPUs, and a rich catalogue of software enables more efficient workflows.
- Storing and archiving big data offline supports research teams working together across projects and enables the most data intensive research workflows.
- Performing advanced data analytics and opening up the world of artificial intelligence to discover new insights and resolve complex problems.
- End-to-end integration supporting high performance data transfers across institutional boundaries allows you to quickly and efficiently transfer big data to and from NeSI.
- Virtual laboratories, which are a customised, integrated and easy to use one stop shop of domain specific tools and data, enabling research communities to work together.

Some of these features were previously available internally at one or other collaborators. With this procurement, NeSI is making these capabilities available nationally. These features will be rolled out over time, alongside training to ensure ease of use.

NeSI's Greg Hall and Fabrice Cantos at the Cray factory checking in on our XC50

> "With NeSI, each time I can use 1,000 cores and that makes a big difference."
>
> TING WANG, DEPARTMENT OF MATHEMATICS AND STATISTICS, UNIVERSITY OF OTAGO

Responding to expectations of emerging data intensive research methods

A strong signal from researchers during the early co-design of the new platform was of expectations NeSI would support emerging data intensive research methods, including artificial intelligence (AI), machine learning, and graph analytics. The new platform recognises the convergence of HPC and big data, providing first class support for high performance and scalable data analytics and AI. This is a new area of capability for NeSI, so the focus is currently on building skills within our own team to ensure we can support researchers as they seek to make use of these new features.

Establishing joint ownership of shared research infrastructure assets

In support of the new infrastructure, NeSI has established a shared ownership model. The goal is to shift behaviours from sole ownership of specific infrastructure assets, to shared ownership of common assets. NeSI is supporting its collaborators to adopt a Tenants in Common shared ownership model, which ensures assets remain owned by institutions able to accept any related risks. Further, this approach enables each investor to realise a mix of benefits from their investment while avoiding the need to duplicate investments. This shared ownership model is underpinned by a single financial model for Total Cost of Ownership, ensuring equality of treatment in accounting for Collaborator Contributions.

> "When you're running several big CFD problems and looking at problems on the order of hundreds of thousands or millions of degrees of freedom, you need HPC."
>
> DAVID LADD, AUCKLAND BIOENGINEERING INSTITUTE, THE UNIVERSITY OF AUCKLAND

Gaining benefits from Cloud computing

After carrying out exploratory work prototyping seamless integration with public clouds during 2016, the team fed key insights from that work and elsewhere into the design of the new infrastructure. A key requirement was the support for the OpenStack cloud platform, which has quickly emerged as the defacto standard globally for what are often known as open research / science clouds. NeSI's new virtual lab and visualisation services will be based on OpenStack, which will also support a variety of other shared services core to the new platform. Having OpenStack as a platform in common with other national advanced research computing infrastructures offers new opportunities for international collaboration and reduces costs to acquiring new capabilities. With the transition platform Kupe in place in late 2017, NeSI / NIWA have already built the first Virtual Labs, based on existing solutions in place to support NIWA's science programmes. Investigations are now underway on how best to apply these capabilities to support genomics, building on relations with the NeCTAR infrastructure in Australia to acquire their Genomics Virtual Lab. As the year closes and awareness increases around these new capabilities, several communities are reaching out to NeSI to discuss how they might plan to make use of these features once the new platform is fully commissioned.

KPI 4

Annually, users indicate that services meet their needs

This KPI is measured by a survey which asked users to rate their satisfaction with NeSI services on a scale from 1 to 5, 1 being 'strongly disagree' and 5 being 'strongly agree'.

DEFINITION

Percentage of users in the NeSI annual survey who agree that NeSI services are meeting their needs.

TARGET (Annual)	ACTUAL
>80% User satisfaction	97%

Of the 121 responses (response rate 66%) year to date, 61 researchers Strongly Agree and 57 Agree that NeSI services meet their needs. Of the other 3 responses that did not agree with the statement, feedback was based around reliability and prioritised access, and also difficulty getting started.

Using HPC for research and career development in Life Sciences.

Tom Kelly – University of Otago

“The opportunity to push myself in a new direction and see it come to fruition is why I pursued a career in science and using NeSI was an excellent way to include transferable skills and personal development in my PhD. “

For many early career scientists, certain skills or experience gained through their projects and collaborations can be game changers in determining their future research path. In the case of Otago University PhD graduate Tom Kelly, one of his career turning points came via an introduction to programming and access to NeSI supercomputers.

"Our research group, supervised by Associate Professor Mik Black at Otago, used NeSI extensively. James Boocock and Ed Hills were doing a summer project for NeSI and the Centre for eResearch in the same office when I joined the group as a biologist and mathematician with zero programming experience," he says. "I was exposed to NeSI from day one and all of my postgraduate studies have been entirely computational."

Over the last four-and-a-half years, Kelly has used statistical approaches to study gene expression and genetic interactions, looking at the consequences of gene regulation for populations evolving and the dysregulation in cancers. Access to high performance computing (HPC) clusters, knowledge of the Linux interface, and use of programming languages such as R, became essential for gathering results for both his Honours and PhD projects.

"Drawing upon my mathematical skills and research interests, I've actively pursued research questions in computational biology that require statistical, computational, and HPC techniques to address," he says. "Performing statistical tests across thousands of genes, and adjusting for multiple comparisons, can sometimes go beyond the scope of what a local machine is capable of, especially when looking at interactions between pairs or combinations of genes."

Whether he's analysing biological pathways or testing a hypothesis where the underlying distribution is unknown, Kelly has come to rely on HPC to help him delve deeper into complex datasets.

"The role of HPC in assembling genomics is well known but there are other ways HPC is used in bioinformatics as well. I usually deal with processed data but the statistical procedures can also be computationally demanding," he says. "It's been particularly fulfilling to gain sufficient scientific computing and HPC skills to be able to participate in eResearch conferences and have the expertise to teach programming workshops to students in the same situation as I was not so long ago."

From time to time, he called on Matt Healey and Peter Maxwell from NeSI's Solutions Team, who were always quick to help with questions both big and small.

"The main thing that I appreciate is that they've been very patient with jobs that didn't work as intended," says Kelly. "I've never hesitated to ask 'dumb' questions and it's been a great system to try things out and develop expertise. I'm increasingly aware that core-hours cost money and it's been a privilege to learn on the institutional access where this has not confined my learning opportunities or the research questions that I could pursue."

With data analysis and computational techniques driving many of today's research approaches, a greater need is emerging for researchers to build skills in these areas.

"We have the technology to generate vast amounts of genetic data and the computing resources to process it but very few people are equipped to handle it effectively," Kelly says. "I support the Software Carpentry initiative wholeheartedly and am glad for the role NeSI has taken with it in New Zealand. I encourage anyone in the life sciences to take any opportunity to learn computational tools and see if it assists your research or suits you as a future career direction."

Kelly recently completed a PhD project that focused on methods for addressing synthetic lethality in breast and stomach cancer. Synthetic lethality occurs when deficiencies - caused either by mutations, inhibitors or other factors - affects the expression of two more more genes and leads to the death of the cell.

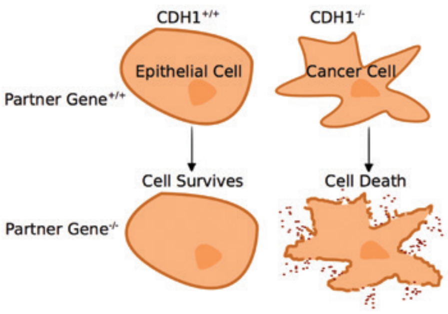

Synthetic lethality in cancer: Rationale of exploiting synthetic lethal interactions for specificity against a tumour suppressor gene (e.g. CDH1) while other cells are spared under the inhibition of a partner gene.

To address this, Kelly explored how to develop a bioinformatics tool to identify synthetic lethal partners of a (cancer) gene from expression data and design indirectly targeting medicines against genes lost in cancers.

His investigations were further extended to analyse and simulate synthetic lethal pathways to examine whether the method could be applied to molecular pathways in cancer. Access to NeSI has been essential yet again.

"These aren't as common in bioinformatics but the simulations, in particular, became a crucial part of my thesis to support the methods that I had already developed, applied, and released as R packages," he says. "The simulations became a very computationally-intensive part of the project and are a lot stronger and more conclusive for the heavy use of HPC. Simulating large gene expression datasets and complex biological pathway networks required HPC to be performed and this aspect of the project was only pursued to this depth because of the skills I had developed in using R, Linux, and NeSI up to this point."

With his PhD thesis now complete, Kelly has begun applying for postdoc positions overseas, with a particular interest in a career in Japan. In Kobe, RIKEN operates the K supercomputer, named after the Japanese word "kei" for the number 1016 and is currently ranked as the eighth-fastest computer in the world. Many research groups at the RIKEN research centres and Universities in Japan use the K supercomputer for their research. A postdoctoral position there would draw upon the programming and HPC experience Kelly gained with NeSI and provide opportunities to take his skills to yet another level.

Reflecting on how he has gotten to this point in his career, Kelly says his experience with HPC has been valuable for more than just the practical skills.

"I think the biggest thing that I have gained is experience tackling new problems with new technologies and having it pay off," he says. "The opportunity to push myself in a new direction and see it come to fruition is why I pursued a career in science and using NeSI was an excellent way to include transferable skills and personal development in my PhD. HPC is definitely a useful resource to have access to and something I would encourage early career researchers to get some experience with. It will open up research questions that you would not have even considered otherwise."

Predictive models to combat invasive species.

Andrey Lustig – Lincoln University

“To run a large number of simulations and scenarios, we needed more computational power.”

A predictive model developed by a Lincoln University researcher using NeSI supercomputing resources could help in New Zealand’s efforts to combat invasive species.

The Bio-Protection Research Centre at Lincoln University is a national Centre of Research Excellence (CoRE) that aims at finding innovative, natural and sustainable solutions to protect New Zealand’s plant-based, productive ecosystems from pests, diseases and weeds.

Working with Associate Professor Susan Worner, Audrey Lustig’s doctoral project at Lincoln used NeSI resources to build a computational model that improved researchers’ ability to track and correlate the interactions between an invasive species and their ecosystem.

What made this model unique is its inclusion of movement and dispersal possibilities to help determine establishment success and spread in at large spatial scales. To date, many of the existing models ecologists use for studying invasive species fail to account for the species re-locating within the space.

“The difficulty that arises with movement is that it adds a layer of complexity to your model because you have more parameters to estimate and to fit into the model,” Lustig says.

Her model linked Individual-Based Models (IBMs) of insect species spread with spatial information gathered through Geographical Information System (GIS) technology.

“Basically it’s a model that allows us to fit, within one framework, all the knowledge we know about the species, including different vectors of dispersal,” she says.

Using this model, researchers and government agencies can determine where and when a species may establish and spread, and then make decisions about what actions need to be taken.

“If we have a clue of how the species acts, we can do different scenarios analysis, by varying its dispersal or reproductive abilities for example, such as we would do for climate projects, to imagine different scenarios for the future,” she says.

Risk mapping tools like this have already helped NZ tackle some of its most persistent pests. For example, in 2010 when the great white butterfly was detected near Nelson and posed a high biosecurity risk to crops of cabbages and broccoli, government agencies used computer modelling to explore and test possible management strategies. By 2016, the great white butterfly was eradicated from New Zealand.

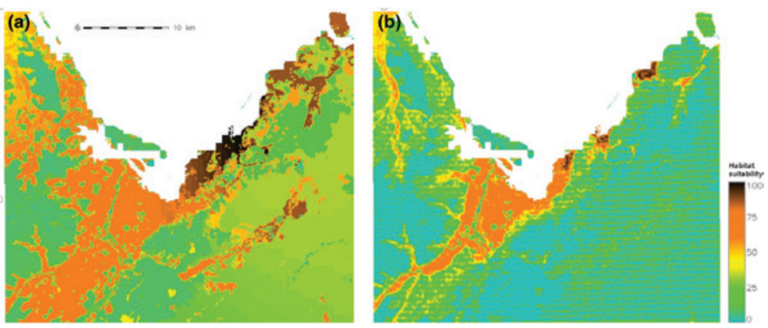

FIGURE 1 Two suitability maps used to build the survival layers of *Pieris brassicae*. (a) Survival layer, LS_1, including the necessary base temperature to complete the *P. brassicae* life cycle within a season (GDD) and the New Zealand Land Cover Dataset LCDB2 (LCC). (b) Survival layer, LS_2, comprised components of the first landscape, as well as a raster layer that captures local elevation variation (LEV)

The challenge that comes with building complex models, however, is that they demand powerful computing resources. This is where NeSI resources became key for Lustig’s work.

“To run a large number of simulations and scenarios, we needed more computational power,” she says. “We used NeSI to run the different scenarios and to test a wide range of parameters. It would have taken months and months to run on my own machine, but I had the results in about two weeks - so that was incredible.”

This was her first time using the systems, so she called upon the support of NeSI’s Computational Science Team.

“My personal experience was awesome,” she says. “I requested help from one of the NeSI team members and Peter Maxwell personally interacted with me for the entire process. He helped me optimize my code, which I’ve never done before, so I learned a lot from a computational point of view. When it came to actually log into the NeSI network he sent me all the instructions and helped me with the process - it was extremely easy for me.”

Now employed by the Geospatial Research Institute at the University of Canterbury, she has been sharing her models and findings with the US Department of Agriculture (USDA), who is planning to use the model for future pest risk assessments.

And while her doctoral work focused specifically on insect assessment, Lustig is applying her modelling expertise to predator scenarios as well. The NZ government and a variety of agencies are working towards making New Zealand predator free by 2050. Lustig is part of a team working on computational models to test different management scenarios.

Computational methods may be the go-to tools for tackling these kinds of large projects, but the next big challenge, she says, is managing the variety of types, formats, and sizes of ecological datasets being collected.

“On the geospatial side of the field, there is a huge effort at the moment to homogenize the data,” Lustig says. “People are gathering more and more data from many places around the country, but many people don’t know they can access infrastructure like NeSI to actually store the data and have help on how to create a management plan for their research.”

Lustig’s modelling work using NeSI resources has been published in Ecology and Evolution and Landscape Ecology.

Objective 5

Enhance national service delivery consistency and performance to position NeSI for growth

Brian Corrie

Solutions Manager

Building an adaptive and responsive organisation

To support the increasing rate of change underway across the sector and within our own organisation, we're seeing early benefits from organisational development work started mid-year on adaptive work practices. NeSI's teams are evolving and adopting agile practices, primarily Kanban and Scrum, with the goal of delivering meaningful outcomes to our customers. A further internal benefit is the clarity of direction this brings to teams, which enhances their opportunities for self-organisation and distribution of leadership. Putting these practices in place now should improve our success during the final stages of commissioning the new infrastructure, while also supporting ongoing evolution of our services into the future.

Taking final steps to move from local to national

Alongside refreshing infrastructure, the team has now broadened our focus onto the next steps for evolving our services. Consolidating our infrastructure to underpin a truly national platform also means we're now addressing some of the last remaining aspects of our previous 'local' operating cultures. We're revising our services to deliver to a single standard nationally. We're combining our site-focused platforms teams into a single platform operations team. We're changing our user support and project allocation practices to operate in a consistent approach nationally.

Enhancing our ability to track benefits and identify challenges

As we implement changes in moving to our new platform, we're lifting the quality of administrative data we collect, to better inform our understanding of who is gaining benefits, and of where we might focus resources. A recent addition to the team is a specialist in data analysis and reporting, whose focus is on improving our ability to answer important questions and tell our stories, through data.

We have reprioritised existing administrative roles to support this shift in focus, which aligns with stakeholder expectations of improved measurement of activities and impact. This will become most evident once we have fully implemented the new platform, as we roll out a suite of changes which broaden the scope and richness of data we collect.

KPI 5 Annual availability of services

The annual KPI of availability measures service delivery consistency and performance. In 2017 NeSI achieved 97% availability.

DEFINITION

Hours annually for which the platform is able to be used to provide the operational service, divided by total hours in a year, less scheduled downtime, expressed as a percentage (The measurement period for this KPI is 3 months).

TARGET (Annual)	ACTUAL (YTD)
≥98% AVAILABILITY	97% AVAILABILITY

NeSI has managed its platforms to deliver a very high level of availability to researchers in 2017 though the ageing nature of the x86 Pan platform is seeing an increasing number of hardware failure issues, as Pan has reached the end of its useful life.

PLATFORM AVAILABILITY

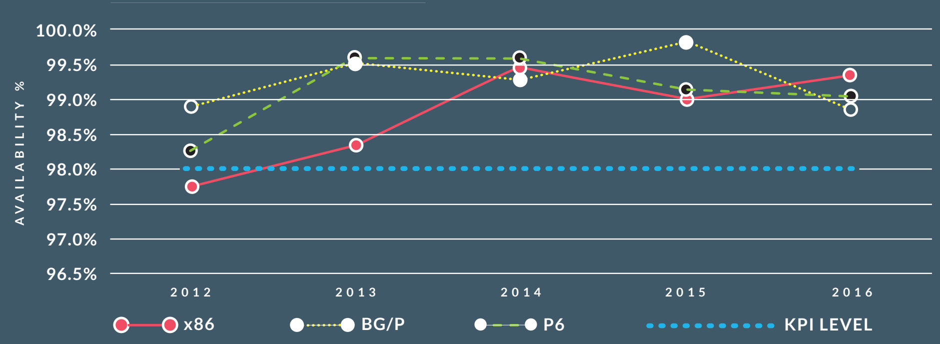

Availability of NeSI platforms, percentage of operational time that the platforms are able to be used.

Objective 6

Realise financial contributions and revenue targets to enhance NeSI's sustainability

Our Collaborators' financial commitment, as a ratio of Crown contributions, is a financial indicator of the health of the NeSI partnership.

We are looking at a change in the calculation of KPI 6 to include accrued Reinvestment Obligation. Current calculation has the Reinvestment Obligation spent only and currently we have not expended all capital on the platform procurement (due to installation phasing) so this appears to be low versus contract. Our forecast after procurement and once we implement the amendment in the treatment of the calculation is to meet the ratio.

Robin Bensley
Operations Manager

Nick Spencer
Site Manager – Manaaki Whenua Landcare Research

David Maclaurin
Site Manager - University of Otago

Fabrice Cantos
Site Manager - NIWA

Marcus Gustafsson
Site Manager - University of Auckland

KPI 6

Contract to date, ratio of collaborator commitments to Crown contribution

DEFINITION
Ratio of Collaborator commitments to Crown contribution from contract start to date (life-to-date).

DEFINITION LTD (Actual)	ACTUAL (Budget)
79.4%	95%

"Through this project, we've already begun to see the benefits of using the NeSI platform for some other projects we would like to get underway in the future."

DAVID BURLACE, SCIENCE & POLICY, DEPARTMENT OF CONSERVATION

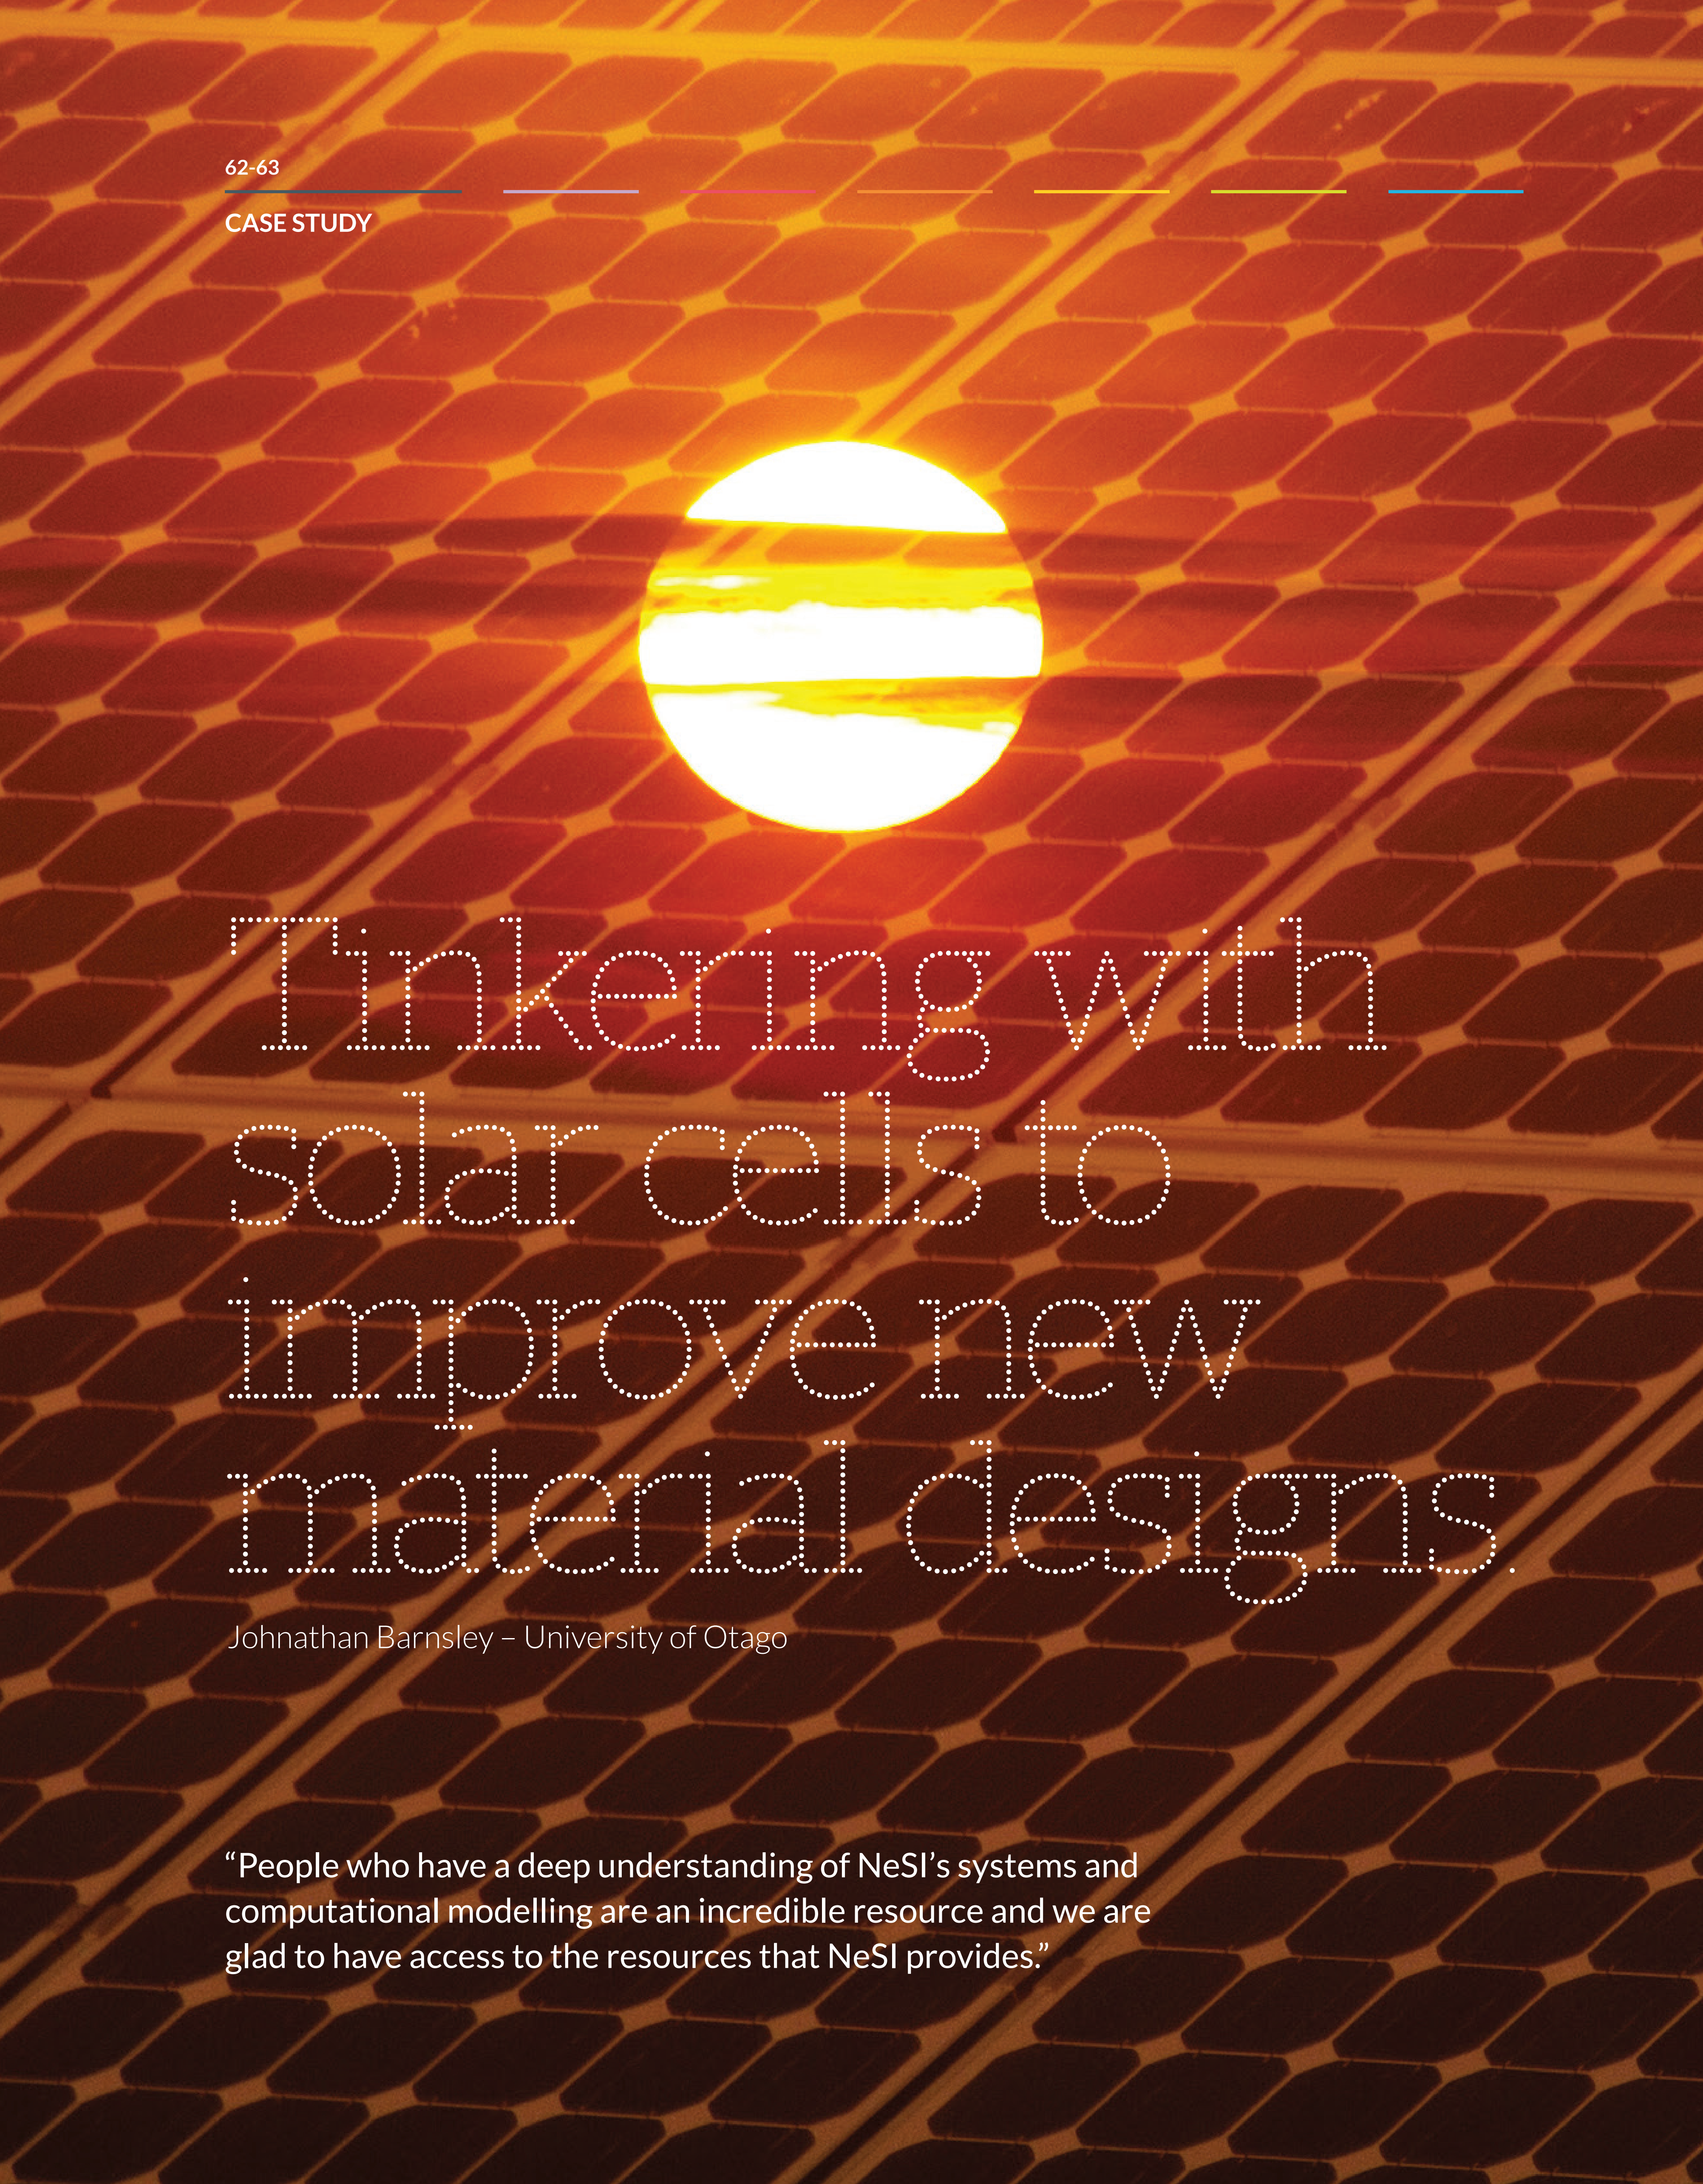

Tinkering with solar cells to improve new material designs.

Johnathan Barnsley – University of Otago

“People who have a deep understanding of NeSI’s systems and computational modelling are an incredible resource and we are glad to have access to the resources that NeSI provides.”

When your research revolves around material properties, you've always got to be on the lookout for interesting behaviour. Jonathan Barnsley knows this well. As a Chemistry PhD candidate at the University of Otago, he studies Donor-Acceptor (D-A) materials. With help from NeSI supercomputers, he's uncovering some interesting findings on a molecular level that could have big impacts across a number of industries.

"It is very exciting to be involved because you're constantly learning," Barnsley says. "It just takes somebody to notice something that you haven't noticed in the past and that's really where things come from. You never know what's around the corner."

D-A materials have an ability to exchange charge from an electron rich donor to an electron acceptor, forming a formally positively charged donor and negatively charged acceptor. Barnsley and his supervisor, University of Otago Chemistry Professor Keith Gordon, are exploring ways to apply the most effective qualities of D-A interactions into the design of new and novel materials.

Some of the materials they work with are D-A dyes, which can be used in dye-sensitized solar cells. Improving the efficiency of solar cells has obvious applications in the green energy sector, but findings from their work could also be applied to other technologies, like sensors, and influence the design of materials in other sectors, such as photodynamic therapy or nonlinear optics. Their latest NeSI-supported work has been published in two papers recently, one of which is featured on the cover of ChemPhotoChem.

When investigating D-A structure-property relationships, one of their first steps is to use techniques under the umbrella of spectroscopy. In simple terms, they shine a laser, or bright light source, at a sample of the material and depending on how the light reacts, either passing through or being scattered, the behaviour provides information about the material's properties. Those experimental results are then compared with data gathered from computational models run on NeSI computing resources and smaller, specialized clusters at the University of Otago.

"That gives us an enriched dataset to really understand how these systems are behaving when they interact with light," Barnsley says. "Leaps of understanding have been achieved, in part due to some of the computational approaches. We find it is very useful to have that modelling approach to understand our systems."

This mix of experimental and theoretical approaches is especially useful when experimental data isn't always obtainable. For example, they can use vibrational spectroscopy to probe a material's structure and then combine that with density functional theory to model the material's electronic and geometric parameters.

"There are some things that computational chemistry can probe that we simply can't observe in the lab, at least at this point in time," Barnsley says. "It would be difficult for any scientist to remain on the frontiers of dye design and characterisation without computational approaches."

The predictive capabilities of computational methods also open the door for better observation and understanding of properties like structure, absorption profiles and transition nature — key characteristics that can impact the performance of a material's design.

"If you're a synthetic chemist, you can spend a long time making a material and find out eventually it doesn't have the properties you were hoping for, whereas computational chemistry is becoming better and better at predicting properties ahead of time," Barnsley says. "That's a tantalizing thing that may come into play for us in the dye design field."

Barnsley and Gordon work with a number of local collaborators, including fellow University of Otago researchers Associate Professor James Crowley and Dr Nigel Lucas, as well as some across the Tasman, partnering with the University of Wollongong's Professor David Officer and Dr Attila Mozer in Australia. Their collaborations even extend overseas, with some projects involving research groups led by Professor Johan Bobacka at Åbo Akademi University in Finland and Dr Grzegorz Lisak at Nanyang Technological University in Singapore.

Currently, a primary source of D-A materials comes from Officer and Mozer, whose Wollongong labs are experimenting with new designs for solar cell systems. Samples are sent to Barnsley and Gordon, who tinker with the systems to identify their properties and explore ways to improve their design.

"We are really lucky to work with people from these groups who are doing all the hard yards making the materials. I basically get to 'play' with the materials once somebody's made them," Barnsley says. "My work tries to understand how these things interact with light and how we can perhaps make them better at absorbing light or try different combinations to absorb light at different parts of visible region."

Meanwhile, he also has colleagues developing supramolecular architectures or other complex computational models that he needs to factor into his studies.

"Synthetic chemists are pushing the envelope making wonderfully large, elegant molecules that have all sorts of interesting properties," he says. "We have to keep up with that and be able to run the modelling and simulations of those big molecules."

Case in point, they recently worked on their largest simulation yet, a model with 286 atoms and 1,283 electrons. The analysis demanded significant computing power, much more than their group's computing resources were providing.

"We couldn't do it on our cluster at all. The jobs were taking six months to complete. Factoring in the number of calculations you had to do to analyse these correctly — it was just too long," Barnsley says. "So NeSI helped a lot in that situation, where we had these novel, complicated and large architectures."

While the models Barnsley works with may be getting increasingly complex, the integration with NeSI resources has been fairly simple.

"Some of the proof of how good the experience was with NeSI is that we didn't really need a lot of help. It was set up in a fairly straightforward manner and it made it really simple to come from our small scale system up to the high performance realm," Barnsley says. "Help from people like Ben Roberts, Gene Soudlenkov, and Peter Maxwell was outstanding. People who have a deep understanding of NeSI's systems and computational modelling are an incredible resource and we are glad to have access to the resources that NeSI provides."

With new project ideas already on the horizon, Barnsley and Gordon's use of NeSI won't be slowing down anytime soon.

"The Keith Gordon group will continue to analyse novel and exciting materials, so for us there will always been a need for quality high performance computing," Barnsley says. "I'm looking forward to a lot of new possibilities and applications of the work we're doing. It's as exciting as it's ever been, so it's good to be a part of it."

A national organisation

NeSI is a collaboration of four institutions coordinating investments in partnership with the Crown. This collaboration is constituted through a legal agreement between the University of Auckland, the National Institute of Water & Atmospheric Research Ltd (NIWA), the University of Otago and Manaaki Whenua Landcare Research, with the University of Auckland as the legal contracting entity with the Crown, commonly referred to as the Host. A set of operating principles defined in the Collaboration and Services Agreements forms the basis for governance, management and operations.

Governance

NeSI's Board is responsible for strategy development, policy, approving major initiatives and investments, as well as monitoring the NeSI risk register. The four collaborator institutions appoint three Directors, alongside an independent Chair and another independent member with expertise in the field representing the research sector at large. All Board members are focused on the interests of NeSI, being the effective delivery of national research infrastructure services.

In 2017, the NeSI Board comprised the following members:

- Rick Christie, Independent Chair
- Prof. Andrew Rohl, Independent, Professor of Computational Science, Curtin University
- Dr Murray Poulter, former Chief Scientist, NIWA
- Stephen Whiteside, Chief Digital Officer, University of Auckland
- Prof. Rob Ballagh, Department of Physics, University of Otago (term ended August 2017)
- Justine Daw, General Manager, Manaaki Whenua Landcare Research (term commenced August 2017)

The NeSI Research Reference Group, established in late 2015, is an advisory group made up of eight research community leaders with strong technical knowledge on the impact of eScience as an enabler of research. The Group has a key role in providing timely advice and input on strategic and policy matters of interest to NeSI, especially those most relevant to researchers.

- Dr Joseph Lane, Chair, Senior Lecturer (Physical & Theoretical Chemistry), University of Waikato
- Dr Nauman Maqbool, Group Leader Knowledge & Analytics, AgResearch
- Prof. Cristin Print, Molecular Medicine & Pathology, University of Auckland
- Prof. Ian Foster, Department of Computer Science, University of Chicago
- Dr Sam Dean, Chief Scientist - Climate, Atmosphere and Hazards, NIWA
- Prof. Blair Blakie, Department of Physics, University of Otago
- Prof. Barbara Chapman, Professor, Institute for Computational Science, Stony Brook University

Rick Christie

Chair, Independent Director

Prof. Andrew Rohl

Independent Director and Professor of Computational Science Curtin University

Dr Murray Poulter

Chief Scientist, Atmosphere, Hazards and Energy NIWA (Retired)

Stephen Whiteside

Chief Digital Officer University of Auckland

Prof. Rob Ballagh

Department of Physics University of Otago

Justine Daw

General Manager, Manaaki Whenua Landcare Research

Robin Bensley
Operations Manager

Kirsten Brown
Operations Coordinator

Mike Ladd
Strategic Projects Manager

Laura Casimiro
Operations Administrator

Organisational health

NeSI has developed its team across 2017 with a number of recruits into newly established positions to bring in new capabilities, and through beginning the implementation of agile methodology:

New Positions:

- In the Directorate we have hired a Communications Manager to help us manage change and sector outreach,
- A Data Analyst also joins the directorate to help us be more responsive and strategic in understanding insights from our user base and platforms.
- In our Support team we have increased capacity with an Application Support specialist at NIWA and introduced a Level 1 Support Analyst based at the University of Auckland.
- We are increasing investment into our software engineering team, with a Product Manager to own the change programs in our services and an additional analyst programmer.

Agile Implementation:

- NeSI is moving from being a traditional waterfall project management organisation to be more adaptive in order to build a better team culture and deliver faster.
- We are now running the applicable teams through Scrum and Kanban processes.

Organisational design

NeSI operates as a virtual organisation (by contract) with teams formed along functional lines.

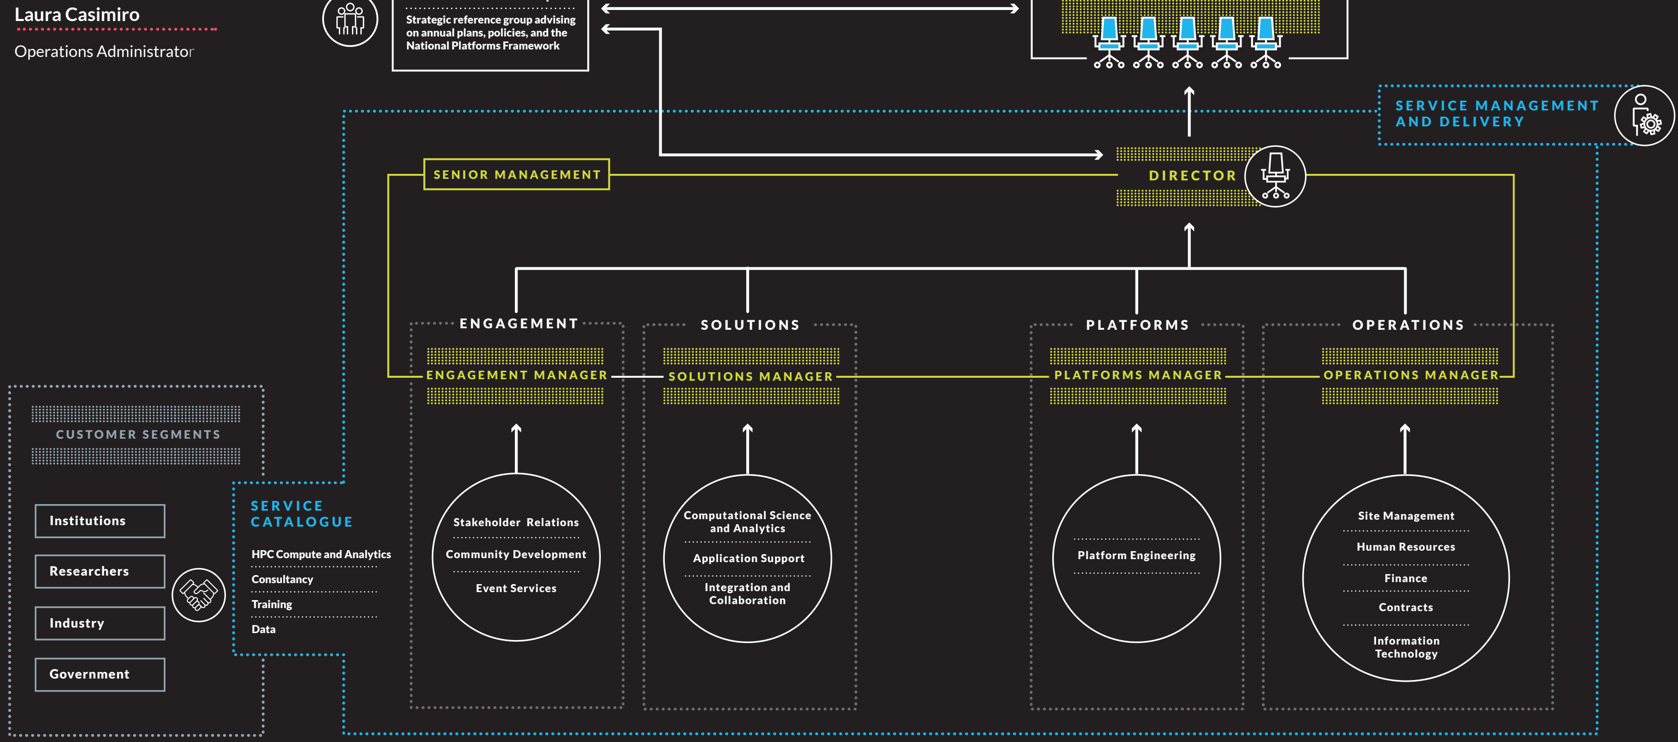

1

A statistical modelling approach to monitoring ecosystems.

New Zealand's Braided River Floodplains are classified as threatened, and preserving their function and biodiversity presents many challenges.

For his PhD, Lincoln University researcher Tyler Brummer set out to respond to some of these challenges. His aim was to gain a greater understanding of the effects of changing river flow patterns, in order to find ways to guide future conservation efforts.

He partnered with Manaaki Whenua Landcare Research, Environment Canterbury and Lincoln University to understand the impacts of river flow changes alien plant species' abilities to invade and alter natural ecosystem processes.

Brummer collected over 2,800 data points from 13 braided river ecosystems in Canterbury, from the plains to the foothills, to rivers fed by the Southern Alps.

He was interested in sites where non-native plant species occurred in multiple river ecosystems. He wanted to determine how "they related to flooding and drought, and whether or not certain species or groups of species were more frequent in abundance in different kinds of rivers."

Brummer was introduced to NeSI's services by a colleague at Landcare Research.

He used NeSI resources to produce complex statistical models and compare the instances of alien plant species with the number of native plant species in each ecosystem. He looked for patterns in how each species interacted with 16 different environmental variables, including river flow, climate, land cover, and human impact.

In the end, Brummer ran over 16,000 different simulations. With NeSI's help, his analysis took a couple of days, rather than several months.

Access to NeSI's supercomputing also meant Brummer could conduct in-depth analysis and increase the complexity of his research, which enabled him to gain a better understanding of the various interacting forces within the ecosystems.

Brummer's insights in his final thesis, "Effects of flow regime on the distribution, richness and abundance of alien plants in braided rivers of New Zealand" are significant. His research suggests that management responses need to be different for alien plant species and native plants because, not only do alien and native species differ in their reactions to ecosystem variables, but the patterns are also associated with completely different environmental drivers.

His findings could be valuable for policy-makers in the contexts of climate change management and resource use patterns.

Brummer acknowledges that there is always more work to be done to understand the dynamics of New Zealand's floodplain ecosystems, and to find effective strategies of reducing the level of invasion and reinvasion of non-native plant species.

Future research projects will inevitably need the support of NeSI's supercomputers again.

"From my perspective, it was extraordinarily beneficial to have access to NeSI as a resource, it really transformed what I thought I would be able to do. The field of statistics at the moment is really on the edge and is advancing all the time."

2

Enhancing Wireless Body Area Networks design and functionality

Over the past decade, advances in electronics, computer science, and wireless technologies have brought wireless body area networks (WBANs) to the forefront of real-time monitoring.

WBANs consist of a number of wearable or implanted wireless sensor nodes that can monitor biological signals such as blood pressure, heartbeat rate, and body temperature. The information collected from sensor notes is communicated to an external system through a special node called the coordinator.

During his PhD, Mohammad Nekooei, from the Victoria University of Wellington, gained valuable insights into WBANs. He studied ways to improve the reliability and performance of Media Access Control (MAC) in WBANs, to reduce delays in communication between the nodes and external systems.

Whilst acknowledging that WBANs can be used in a variety of sectors, Nekooei particularly focused on their healthcare applications. The reliability of WBANs is especially important in those cases because it is responsible for carrying critical medical data and can signal emergencies such as vital organ failures.

In order to carry out the numerous simulations required for his studies, Nekooei needed a high-performance computing resource that was fast and reliable.

Before using NeSI, each run of his model was taking more than 200 hrs to complete. With his PhD thesis deadline looming, Nekooei needed to do something differently.

"The NeSI technical staff really helped me," he said. "They responded to all of my requests."

By moving to NeSI's computing resources, he was able to get the answers and insights he needed quickly.

"If I hadn't used NeSI, honestly, I really doubt I would have finished my PhD," he says. "Access to high-performance computing helps me get more results and improve the quality of those results and my research."

As part of his work, Nekooei explored ways to improve and automate the design of fuzzy logic controllers (FLCs), which are used in WBANs to handle the flow of information from the sensor nodes. His latest findings were published in Applied Soft Computing in July 2017, in a paper entitled, 'Automatic design of fuzzy logic controllers for medium access control in wireless body area networks - an evolutionary approach'.

With his PhD now complete and further journal papers in the works, Nekooei says he is keeping NeSI in mind as he prepares to enter into the next chapter of his research.

"If I want to continue my research, I will definitely still need to use high-performance computing."

> "High-performance computing is essential for projects like mine when you need reliable and fast platforms to run these jobs."

3

Improving earthquake forecasting methods

Accurate advance warnings of large earthquakes are vital to activate early response efforts and save lives.

Unfortunately, there's still uncertainty about when the next large earthquake will hit.

With the aim to find solutions to this problem, Dr Ting Wang, a senior lecturer at the University of Otago, builds and tests statistical models. The hope is that, in the future, these models will better predict when a large earthquake may strike.

Wang uses earthquake data sequences from Northern California and the Middle America Trench. Many large earthquakes have occurred in this area, and the data is available online for researchers anywhere in the world to use and analyse.

Using this data, Wang wrote a program for running simulations, but she found it too slow to return the results she needed.

Wang needed help from someone with more programming expertise. She decided to enlist the help of NeSI's Solutions Team for two projects related to her earthquake research.

The aim of her first project, funded by the Earthquake Commission (EQC), was to develop a statistical model that automatically classifies earthquake activities into distinct phases, such as foreshocks, mainshocks, and aftershocks.

Not only did Wang receive the programming support she was after, she also accessed NeSI's Pan supercomputer. This meant she could run calculations using multiple scenarios and get results quickly, in just a matter of days.

Wang was supported by Peter Maxwell, a member of NeSI's Solutions Team.

"Whenever I had a problem I would just email Peter and he would immediately reply and help me with it."

Since then, and with the help of another NeSI consultant, Wolfgang Hayek, Wang's programs have been refined even further.

"[Wolfgang] actually looked through my program line by line - and it's a really long program - and he made things more efficient," Wang says. "It's even better than before. I'm testing it extensively and it's working really well."

Wang's second project, funded through a Marsden Fund Fast-Start Grant, studies how earthquake and volcanic eruption data vary over time.

The main challenge from this project was the missing data points because, historically, people weren't as diligent or as capable of recording geophysical activity.

Wang and her PhD student Amina Shahzadi investigated ways that missing data points can be accounted for and factored into sequence analysis and hazard forecasts.

NeSI's support and expertise were needed again. Under the guidance of NeSI's Alexander Pletzer, Wang and Shahzadi were able to run simulations using a larger data set than they'd worked with previously and have results to study in hours instead of days.

"If we want to solve real-world problems, we need sophisticated models to capture what's happening in the real world, which is not simple."

4

Kiwi chemists develop new analysis tool

A better understanding of molecules, especially their structures and properties, leads to new products and applications in a variety of industries, including the development of industrial coating, medications, fuels and agriculture.

The properties of a molecule are dictated by its structure. Structures can be transformed through molecular conformation, whereby the atoms of a molecule adopt and convert between different configurations, especially by rotation about individual single bonds. Even a small change in a molecule's structure has a big impact on its properties, such as its shape, function and interactions with other molecules.

Using computational modelling, scientists can experiment with conformer ensembles – sets of conformations – in order to further understand a molecule's structure.

University of Canterbury researcher, Nathaniel Gunby, wanted a program that could generate and analyse conformer ensembles, with an aim to gain more insights on non-organic molecule structures.

With his supervisors, Dr Sarah Masters and Dr Deborah Crittenden, and using a NeSI supercomputer, Gunby developed the Universal Conformer Generation and Analysis Program package (UNCONGA).

A versatile computer program, UCONGA can generate a variety of conformer ensembles, covering a wide gamut of molecular elements and ring structures.

The program's functions include analysing sets of conformers, gathering similar conformers together so that representative or unusual conformers can be chosen for further study, and sifting out sets of conformers that are too similar to each other.

Researchers using UCONGA can create improved models of non-organic molecular structures, which enables them to better predict both molecular properties, and the rate at which reactions will take place. It's become a particularly useful tool for computational and structural chemists.

Without NeSI's resources, Masters says her group would have struggled to complete the development of UCONGA in time, and Gunby would have found it difficult to meet his thesis deadline. Using NeSI, they could also easily identify and fix bugs in the program.

"Without a doubt, we would not have been able to undertake the development work for this project as quickly or efficiently without the support from NeSI." Masters says.

Looking ahead, Masters and her group are exploring ways to improve industrial coating processes to reduce the effects of wear and tear on machine parts. UCONGA will be one of the analysis tools they use to support those investigations.

> "Without a doubt, we would not have been able to undertake the development work for this project as quickly or efficiently without the support from NeSI."

5

Native forests absorbing more carbon dioxide

With the devastating effects of climate change seen globally, New Zealand is environmentally and politically committed to find the best ways of mitigating some of these effects.

Research has shown that carbon dioxide is responsible for most of the human-induced atmospheric warming. Fortunately, we also know that oceans and forests lessen the damaging impact by absorbing carbon dioxide from the atmosphere. These natural assets that uptake carbon are known as carbon sinks.

The good news is that new research led by NIWA scientists Dr Kay Steinkamp and Dr Sara Mikaloff-Fletcher, indicates that New Zealand's native forests (of which there are over 6.2 million hectares) absorb up to 60% more than previously thought. This figure is compared with the Ministry of Education's annual reporting on New Zealand's carbon uptake by all New Zealand forests for the same period (2011-2013).

Mikaloff-Fletcher and Steinkamp's study is surprising not only because of the significant carbon uptake from New Zealand's native forests, but because of what they discovered about where and how the highest uptake happens.

Typically, high levels of carbon uptake are associated with newly planted forests, not relatively undisturbed mature forest, explains Mikaloff-Fletcher.

For their study, they needed a modelling approach that could estimate the amount of carbon uptake from different sites around New Zealand, both on land and at sea.

Using NeSI supercomputers, they adopted an "inverse" modelling approach to estimate the level of carbon uptake through measurements of carbon dioxide present in the atmosphere.

They gathered information from a network of New Zealand sites, including NIWA's clean air station at Baring Head near Wellington, its atmospheric research station at Lauder in Central Otago, and observations taken from a ship that travels between Nelson and Osaka, Japan.

In order to connect observed changes in the atmosphere with carbon uptake, they needed detailed information about where the air came from before reaching the site. The team therefore used a high-resolution weather forecast model run on NIWA's FitzRoy supercomputer to simulate the winds over a region and pinpoint the uptake areas.

"The story the atmosphere is telling us is that there's a big carbon sink somewhere in the South Island, and the areas that seem to be responsible are those largely dominated by indigenous forests," says Steinkamp. "However, we cannot rule out an important role for carbon uptake in the hill country or from pasture from our current data."

Mikaloff-Fletcher adds that NeSI's computing resources were integral to their project.

"New Zealand's topography has a major impact on the airflows over our region, so high resolution modelling is essential to be able to use this technique effectively in New Zealand."

"High resolution modelling is essential to be able to use this technique effectively in New Zealand. NeSI support opened the door for us to work with models running at the spatial resolution we need."

6

NeSI partners with Software Carpentry to expand research computing training

NeSI has been a Software Carpentry Foundation member since 2015. This year, it ramped up its collaboration level to further drive the development of local training communities in New Zealand.

In June 2017, NeSI signed a Platinum partnership to increase the number of Software Carpentry and Data Carpentry instructors trained in New Zealand, support the delivery of more workshops across the country, and ultimately strengthen and build local training communities.

As part of its core curriculum, Software Carpentry delivers a two-day workshop. Topics covered include how to automate repetitive tasks, using version control to track and share work, and building programs that can be read, re-used and validated in programming languages. Its sibling project, Data Carpentry, focuses on teaching data analysis skills.

Both courses target graduate students, early career faculty, and "anyone who is accumulating data faster than their training or their community's training can deal with," says Jonah Duckles, Executive Director of the Software Carpentry Foundation.

He adds, "NeSI is a catalyst for institutions to discover the ways that Software Carpentry can be broadly beneficial."

Duckles sees the potential for New Zealand to act as a testbed for refining and expanding Software Carpentry's offerings. For example, with NeSI's connections to many different universities, there's scope to improve the specialised courses Software Carpentry offers, such as Library Carpentry. Duckles wants libraries to think about the role they can play in upskilling and training researchers.

By engaging with campuses, NeSI could get libraries "excited about how they could be a gateway, in conjunction with other research computing groups, to this skills-building and training. That makes NeSI's value proposition of resources, people and infrastructure really valuable in the long-term."

Several NeSI staff are already trained to deliver workshops and have been doing so since 2015, however, to keep pace with the growing demand for workshops, more instructors are needed. NZ needs more people who can train the trainers.

"These next few years will be important as NeSI identifies how to continue building momentum in this space – whether it's through training more instructors, helping our members and other organisations host workshops, or encouraging our members to become Software Carpentry members as well," says NeSI's director, Nick Jones. "We're confident these activities will advance computing capabilities across New Zealand's research system, and we're excited to see how NeSI and New Zealand researchers can contribute to the development of Carpentry's global communities."

7

Using GPUs to expand our understanding of the Solar System

Dr. Philip Sharp, a senior lecturer at the University of Auckland, is interested in the early evolution of the Solar System.

Sharp collaborates with Professor William Newman, of Physics & Astronomy at UCLA. Together they investigate whether Nice, the leading model for the Solar System's evolution, accurately predicts the current orbits of the four giant planets, Jupiter, Saturn, Uranus and Neptune, as well as the Sun and minute planets called planetesimals.

In his research, Sharp uses mathematical methods called number of body (N) simulations.

"Collisionless N-body simulations over tens of millions of years are an important tool in understanding the early evolution of planetary systems." Sharp says.

Using a single Graphics Processing Unit (GPU), Sharp and Newman ran simulations between 20 and 100 million years with N=1024 astronomical bodies.

Other researchers have used this N-value before, but Sharp and Newman are the first to perform accurate simulations with it. Their findings were published in the Journal of Computational Science.

Based on their research, Sharp and his colleagues argued that the Nice model was unstable for the number of bodies researchers use, therefore limiting Nice's predictive power.

To strengthen their argument, they sought to perform accurate simulations with larger values of N. However they were limited by their existing simulation program, because the elapsed time would be too great.

Chris Scott from NeSI's Computational Science Team came on-board and helped modify the program to run on multiple GPUs simultaneously. With Scott's help, the researchers were able to convert their simulation program to make use of multiple GPUs and apply it to a range of numbers of bodies, from 1,024 to 65,532.

They immediately observed improved performance, particularly with the larger numbers of bodies where the observed speed increased by 1.7 for simulations with 16,384 or more bodies.

Next, Scott experimented with OpenMP, an Application Programming Interface (API) commonly used for parallel programming. The combined speedups from using two GPUs and OpenMP were 1.9x for 4,096 bodies, 2.0x for 16,384 bodies, and 1.8x for 32,768 bodies.

"These gains will allow researchers to run simulations of larger systems in a more reasonable time and verify the results obtained from their calculations," says Scott.

Sharp is very reliant on high performance computing, and is grateful for NeSI's support and resources, saying they are the best in New Zealand.

"Put simply, the work would not be possible without HPC. We chose NeSI because its HPC resources are the best in the country."

Water conservation applies to weather models too

The National Institute of Water and Atmosphere (NIWA) has a Weather Related Hazards research programme, which aims to improve New Zealand's ability to accurately predict severe weather and its consequences, such as flooding and storm surges.

A critical part of this research is the ability to use the information from rainfall forecasts to accurately predict the possibility of a flood event.

To carry out this work, NIWA uses a weather model called the New Zealand Corrective Scale Model (NZCSM). Using high-resolution forecasting techniques, NZCSM can predict rainfall levels throughout New Zealand. NZCSM's output data is then transferred into a hydrological system called TopNet.

In order to keep flood forecasts accurate, it's critical to maintain the integrity of the data when transferring between the systems, so as not to lose or gain any water data.

But the challenge of coupling these weather and hydrological systems is that they each use different grid models.

Therefore, the source data from NZCSM needed to be re-mapped before being incorporated into the target grids in TopNet.

To tackle this challenge, Alex Pletzer and Chris Scott from NeSI's Solutions Team worked with Dr Céline Cattoën-Gilbert, a NIWA Hydrological Forecasting Scientist, and Dr Trevor Carey-Smith, a NIWA Climate Scientist.

The solution involved using software that Pletzer and Scott had previously developed as part of a Deep South National Challenge project, and applying it to this data remapping task.

The tool developed by Pletzer and Scott understands how different types of grids store their data and the underlying assumptions behind them, such as if the grids are longitude-latitude based. A researcher simply needs to provide the source data and target grid, and the software will figure out what to do.

"It was great to be able to leverage the outcome of a previous NeSI consultancy project," says Carey-Smith. "After an initial meeting, Alex and Chris provided a robust solution that runs faster than expected and is easy to use. NeSI's support on this project has been extremely valuable and will improve the accuracy of our flood forecasts."

Cattoën-Gilbert's feedback was equally positive.

"It was fantastic to have Alex and Chris's help and expertise. We have found that NeSI's consultancy service has enabled us to make big strides in our research projects."

"NeSI's support on this project has been extremely valuable and will improve the accuracy of our flood forecasts."

NOTES

NeSI
New Zealand eScience
Infrastructure

nesi.org.nz