



NeSI
New Zealand eScience
Infrastructure

National Platforms Framework

2015 Revision

Table of Contents

Recommendations	3
Approval by NeSI Board	3
National Platforms Framework (2015 revision).....	4
Review of Current Platforms and Usage.....	5
Platform Features	5
Platform usage.....	6
Research Needs Survey Analysis.....	6
Summary of survey responses by research domain	6
Key insights from the survey.....	7
Implications for the evolution of NeSI's services.....	7
Data Services.....	7
HPC Compute and Analytics.....	7
Consultancy and Training.....	8
Review of technology directions.....	9
HPC Procurement.....	9
NeSI HPC platforms requirements	9
Desirable Requirements.....	9
Implications for infrastructure architecture	10

Recommendations

The National Platforms Framework outlines NeSI's plan for investment into platform assets and services which meet sector-wide research needs. The 2015 review confirmed the following key aspects of the Framework:

- Consolidation from three to two HPC platforms nationally by mid-2016,
- Replacement of these two platforms along with any platform requirements to meet national Genomics needs by mid-2017 through an integrated procurement process,
- Optimisation across both national platforms to improve operating efficiency through fit for purpose use, including extension through a Cloud-burst model to meet peak demand,
- Extend scope of team expertise to incorporate Data Analytics expertise to support anticipated broadening of researcher's requirements.

This approach provides an opportunity to deliver an infrastructure which will:

- 1) Significantly reduce the barriers to users moving between the Capacity and Capability Platforms, and between Genomics and HPC platforms,
- 2) Provide users with a common development and job management environment, as well as advanced HPC, Data Analytics, and Genomics capabilities, and
- 3) Reduce the cost of Platforms support services while providing improved resilience and support functions.

The revised National Platforms Framework has been informed by consultations with NeSI's Collaborators and Subscribers, and with NZGL, and analyses of:

- 1) Current NeSI Platform utilisation,
- 2) The expected research needs of those users who responded to a research needs survey, and,
- 3) Anticipated HPC technology trends and roadmaps,

The revised National Platforms Framework will support all six NeSI Objectives:

- 1) Support New Zealand's research priorities.
- 2) Grow advanced skills that can apply high-tech capabilities to challenging research questions.
- 3) Increase fit-for-purpose use of national research infrastructure.
- 4) Make fit-for-purpose investments aligned with sector needs.
- 5) Enhance national service delivery consistency and performance to position NeSI for growth.
- 6) Realise financial contributions and revenue targets to enhance NeSI's sustainability

Approval by NeSI Board

The 2015 review process was completed with the NeSI Board approving the 2015 revision at their meeting of March 14, 2016.

National Platforms Framework (2015 revision)

The National Platforms Framework 2015 revisions is as follows:

Table 1: 2015 Revision of the National Platforms Framework.

National Platforms Framework (2015 revision)	
2016	<ol style="list-style-type: none"> 1) Use capacity planning to determine requirements for operational Cloud-burst service by 30 June 2016 2) Decommission University of Canterbury Platforms by 30 June 2016 3) Optimize and sustain fit for purpose use of the existing infrastructure 4) Recruit a Data Analytics expert by 30 June 2016 5) Agree Data services strategy and feed into platform design by 30 June 2016 6) Design platform solutions that will enable NeSI to meet its Goals and Objectives, and develop Requests for Proposals for both Capacity and Capability Systems by 15 July 2016 7) Issue Request for Proposals by 30 July 2016: <ol style="list-style-type: none"> a. Initial responses due 15 October 2016 b. Best and final offers due 15 November 2016 c. Select successful vendor(s) by 30 December 2016
2017	<ol style="list-style-type: none"> 1) Decommission Pan and FitzRoy by 30 July 2017 2) Contracting, installation, acceptance testing, in production by 30 June 2017 with user training in July 2017 1) Optimize and ensure fit-for-purpose use of the new Platforms 2) Optimize Data services
2018	<ol style="list-style-type: none"> 1) Optimize and sustain fit for purpose use of the existing infrastructure 2) Review platform investments to inform future investment plans

Review of Current Platforms and Usage

Platform Features

The primary purpose of HPC is to achieve “shortest time to solution”. To this end, NeSI invested in, and operates two classes of HPC Platforms: Capacity and Capability¹, whose features and application domain characteristics are described in Table 1 and on [NeSI’s website](#).

Table 2: General features of Capability and Capacity HPC systems and of the characteristics of research applications that are executed on them.

HPC Class	Platform	Features and Application Domains
Capability	<p>Foster (Canterbury BlueGene/P)</p> <p>FitzRoy (NIWA P575/ P6)</p>	<p>Application Domains</p> <ul style="list-style-type: none"> • Large, highly coupled problems, which have high inter-processor / low latency communications requirements, and typically, very high I/O demands • Ensembles of large highly coupled problems • Less tightly coupled, but still critically dependent on interconnect performance to achieve scalability • Tightly coupled problems that exhibit poor scaling properties require high performance processors <p>Need high level of Reliability, Availability and Serviceability (RAS) – if being used to support decision making in advance of / during national emergencies, high availability SLA.</p> <p>Design Features</p> <ul style="list-style-type: none"> • Highest performance interconnect fabric • Large processor counts on each node • Modest amounts of memory per processor • Highest performance processors (the current exception being the BlueGene architecture, which is no longer available) • Extreme I/O performance and large data storage infrastructure
Capacity	<p>Pan (Auckland iDataPlex)</p>	<p>Application Domains</p> <ul style="list-style-type: none"> • Problems that have low inter-processor communication requirements i.e. are loosely / not coupled • Can utilise thousands of cores, with near perfect scaling (i.e. Embarrassingly Parallel problems) • Problems that can only run on a single core • Problems that can utilise non-traditional/emerging HPC architectures (e.g. accelerators) • Problems that could be executed on a Cloud Service <p>Design Features</p> <ul style="list-style-type: none"> • A low performance interconnect fabric will not significantly impact time to solution • Modest processor counts per node • Large amounts of memory per processor, but a subset may have very large memory (O(4TB per node)) • Highest number of processors affordable (the focus is not on processor performance) • Good, but not extreme I/O performance

¹ These are the standard terms used within the International HPC community. Capability platforms are also known as “Supercomputers”, and Capacity platforms as “High Throughput Computers”.

Platform usage

Fit for purpose use of the installed NeSI HPC platforms has focused on application domains which make most efficient use of the key features of each architecture as specified in Table 1. In particular:

- Pan – nearly all jobs run on one node or less, or are loosely coupled (i.e. Embarrassingly Parallel) and there is high demand for throughput – hence the primary investment in Pan has been on adding more and more processors.
- FitzRoy – nearly all jobs are tightly coupled, run on multiple nodes, achieve good scalability and have very large I/O demands. It is widely used for research that demand high performance processors and interconnect.
- Foster – as a more specialised platform, it has slow processors with a (relatively) high performance interconnect, making it suitable for application to a number of problem classes that scale well on this architecture. It cannot efficiently execute “single” core jobs.

General utilisation of the Platforms follows naturally – with the Capacity systems executing large numbers of small jobs from a large user base ensuring relatively high levels of utilisation, while Capability systems execute a small number of large to very large jobs from a much smaller user base.

Accordingly, Capability system utilisation is much more susceptible to changes in research activity (e.g. it may drop after a long series of integrations comes to an end, and the research team moves to data analysis, and paper writing activities, or a PhD student comes to the end of their research, before the next student ramps up activity.)

Research Needs Survey Analysis

Summary of survey responses by research domain

Key researchers completed a “Research Needs Survey”, with 46 responses received. Table 3 classifies these responses by primary research domain (in a number of cases, the respondents indicated more than one domain), and by the primary platform used by the respondents.

Table 3: Response to the survey by research domain, and by platform used (grouped by science domain).

ID	Primary Research Domain	Responses	Platform Used		
			Not all respondents are NeSI users		
			Pan	FitzRoy	Foster
1	Biomedical Sciences	5	3		2
2	Cellular, Molecular and Physiological Biology	1			
3	Earth Sciences and Astronomy	19	4	13	1
4	Ecology, Evolution and Behaviour	6	4	1	
5	Economics and Human and Behavioural Sciences				
6	Engineering and Interdisciplinary Sciences	5	5		
7	Humanities				
8	Mathematical and Information Sciences	3	1		
9	Physics, Chemistry and Biochemistry	7	7		(1)
10	Social Sciences				
	Total	46	24	14	4

Key insights from the survey

The survey responses covered in Table 3 suggest that:

- 1) NeSI's Capacity platform Pan is operating as planned by supporting research across a wide range of research domains. Further analysis of the detailed responses and of historic utilisation on Pan indicates the large majority of research jobs meet the application expectations of a Capacity class system as noted in Table 1
- 2) There were 46 responses to the Survey from a broad range of research groups. Further end user engagement is anticipated during the National Platforms Framework 2016 review
- 3) The (large) majority of respondents have few or no identified international collaborations / peers in their areas of research. This is of interest as it is increasingly complicated and challenging for individual researchers and small research groups to maintain and develop HPC software codes

Implications for the evolution of NeSI's services

The following points capture key requirements for NeSI's services which were identified by respondents:

Data Services

- a) Faster methods to transfer large datasets between research groups, onto the NeSI HPC Platforms, and to / from international peers / collaborators
- b) Improved data share services
- c) Access to, and management of large datasets (i.e. capability to host reference datasets for long periods)

HPC Compute and Analytics

- a) HPC Compute:
 - i) The big Earth Sciences and Astronomy researchers have a well-articulated understanding of future needs, e.g.
 - (1) For high performance cores (not much use yet for GPGPUs or MIC architectures);
 - (2) Very large core-hour requirements (O(100M core-hour) per annum in the case one specific community;
 - (3) High interconnect fabric performance to enable scalability on tightly coupled codes
 - (4) Large data output and storage (O (1PB) per simulation and the need for multiple simulations
 - ii) Researchers in Biomedical Sciences will also need access to large Capability Platform resources;
 - iii) In some science domains – there are major gains to be made by transitioning to codes that can make use of GPGPUs (e.g. Molecular Dynamics codes such as AMBER) – leading to very cost effective HPC services and improved time to solution metrics.
 - iv) Use of MIC architectures (e.g. the new Knights Landing self-hosting Many Integrated Core architecture) in science codes that deliver performance improvements in time to solution is less clear
- b) HPC Data Analytics:
 - i) The need for Data Analytics, and reduced movement of data (i.e. analytics *in situ*) will be an area of growth in the coming years
 - ii) In part this will be driven by the needs to analyse PB scale datasets
- c) HPC platform operations:
 - i) There is an expressed desire for better access to fit-for-purpose platforms and management of queues / workloads on the HPC platforms.
 - ii) A number of research groups are using NeSI platforms that are not optimal (i.e. fit for purpose to their needs)

Consultancy and Training

- a) Many researchers believe that they could be more effective if they had dedicated scientific programmer resources to draw on
- b) Internationally, there is substantial effort being applied to the issues facing a number of science domains, e.g.
 - i) In genetics research software codes are being developed that remove some of the bottlenecks imposed by current serial codes (typified by users requesting nodes with larger and larger amounts of memory – to be operated on by one core). New Zealand research groups need to begin to adopt these new methods so that they can achieve faster times to solution, and NeSI can make cost effective investments in HPC platforms
 - ii) In Molecular Dynamics, Materials Science and Computational Chemistry, GPGPU acceleration of some codes is already showing big gains in time to solution as well as better benefit to cost metrics

However, researchers are typically conservative, not prioritising the time to test new methods and approaches to problem solving. NeSI's Scientific Programmers can assist researchers to make these transitions, which is likely to lead to major benefits for New Zealand science

- c) With the increasing need for Data Analytics in the future (including visualisation), NeSI will invest in recruiting "Data Analytics" expertise

Review of technology directions

While HPC technology is always evolving, a number of potentially disruptive technologies have come into view over the past year, while others that are new now, will soon reach levels of maturity appropriate for a national HPC infrastructure. The detailed analysis focused on:

- Processor architectures;
- Interconnect developments;
- The development of deeper memory hierarchies;
- Storage systems, that can deliver both high IOPS and bandwidth ;
- Parallel Filesystems;
- High density packaging, power efficiency and cooling;
- Software developments; and
- Cloud options.

Taking all these points together, the new technologies that will become available in 2017 make this a good a time to acquire new platforms, since:

- Storage technologies that are available today, but immature, will be much improved;
- Deep memory hierarchies will be available and will have gained some maturity, as well as the software systems needed to utilise them;
- More efficient and performant processors will be available; as well as
- New interconnect technologies.

HPC Procurement

NeSI's team reviewed international best practice for HPC procurement, noting best practice recommendations were consistent with the recently completed NeSI High Performance Computing Procurement Manual.

NeSI HPC platforms requirements

Desirable Requirements

The key business requirements to be delivered through implementation of the National Platforms Framework include:

- 1) Making it easy for users to develop and run research workloads/jobs and apply HPC compute and Data Analytics tools on either/both platforms
- 2) Fit-for-Purpose platforms that meet researcher needs – Capacity (including high throughput) and Capability
- 3) Access to standard “big data” Data Analytics tools
- 4) High level of interoperability/commonality of (Systems) management and monitoring systems on the platforms
- 5) High IOPS and bandwidth for input/output operations
- 6) High reliability and availability
- 7) Transparent management of data on tiers (from Flash to disk to tape)
- 8) Fastest time to solution
- 9) Minimising the Total Cost of Ownership
- 10) Reducing diversity between platforms

Implications for infrastructure architecture

The following features are anticipated from NeSI's renewed platforms infrastructure:

- 1) Single sign on (with home institution credentials)
- 2) Uniform namespace for home filesystems – or failing that, federated namespace
- 3) Same development environment: compilers, linkers, development tools (e.g. profilers, debuggers)
- 4) Same “core” software – e.g. Data Analytics tools
- 5) Same (or very similar) Systems environment on both platforms
- 6) Common implementation of monitoring tools
- 7) Transparent data movement (between platforms)
- 8) Enhanced data management and archiving facilities
- 9) Supported, highly reliable national-scale platforms.
- 10) HPC Compute and Analytic, and Data Services, that are comparable in scale (per user) and quality to those available to international collaborators.