The NeSI National Platforms Framework

Michael Uddstrom[1], Nick Jones[2] , Nic Mair[3]

[1]NeSI Platforms Manager, [2]NeSI Director, [3]NZGL BioIT Manager

New Zealand eScience Infrastructure

Outline

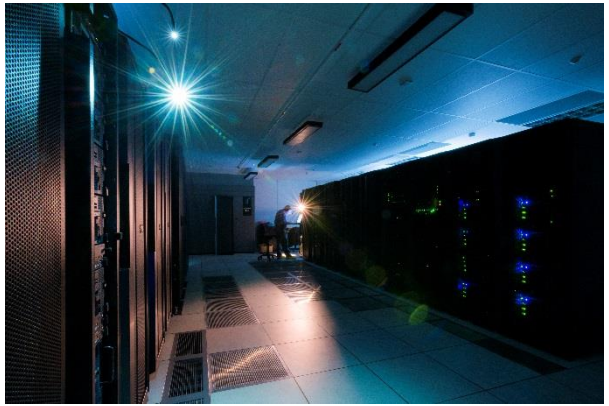

- NeSI.2 on one slide;
- The National Platforms Framework – What & Why;
- The Review Process;
- NeSI Platforms: Current Status:
  - Two definitions – to avoid confusion…
- Current usage – a high level view;
- Planning for the future – Researcher & Stakeholder inputs:
  - Some headlines.
- The Draft 2015 National Platforms Framework on one slide.

Context NeSI.2

- New structure: National management, Service Lines, Service Catalogue;
- National governance over platform investments (CapEx + OpEx)– the National Platforms Framework (NPF);
- New Access Policy covering all allocations:
  - Merit Project grants receive free access to NeSI HPC Compute & Analytics Services;
  - Investor/Collaborator Projects may run on any NeSI Platform (Optimise fit-for-purpose);
  - Allocation model covers HPC core-hours and Computational Science Team staff.
- Service Catalogue:
  - HPC Compute and Analytics, which delivers platform services to researchers;
  - Consultancy and Training (includes Computational Science Team, Applications Engineers) ;
  - Data Services (share and fast transfer).
- 60% of the HPC resource is reserved for Investors, 40% for "Merit" access.

Context: National Platforms Framework Review

- This Framework is the high level plan underpinning "**nationally coordinated procurement**";
- Two relevant Goals to inform the NPF Review:
  - Making it easier to start: Empowering researchers to make effective use of advanced computing capabilities, and
  - Improving time to solution: Enhancing the capabilities available to researchers, and enabling them to address the most challenging problems.
- Four key NeSI Objectives:
  - Support New Zealand's research priorities (NSSI);
  - Increase fit-for-purpose use of national research infrastructure;
  - Make fit-for-purpose investments aligned with sector needs;
  - Enhance national service delivery consistency and performance to position NeSI for growth;

National Platforms Framework

- The National Platforms Framework Review is informed by:
  - Analysis of the usage of the existing platforms (science applications, job size, I/O etc.)
  - Responses to a Research User Survey, which considers:
    - » Current usage;
    - » Planned usage (~3 year time frame);
    - » International collaborations.
  - HPC technology roadmaps;
  - Financial implications for the NeSI platforms investment Fund;
  - Stakeholder input and feedback;
- Is reviewed annually;
- Is approved by the NeSI Board Approval;
- Is implemented by the Platforms Manager.

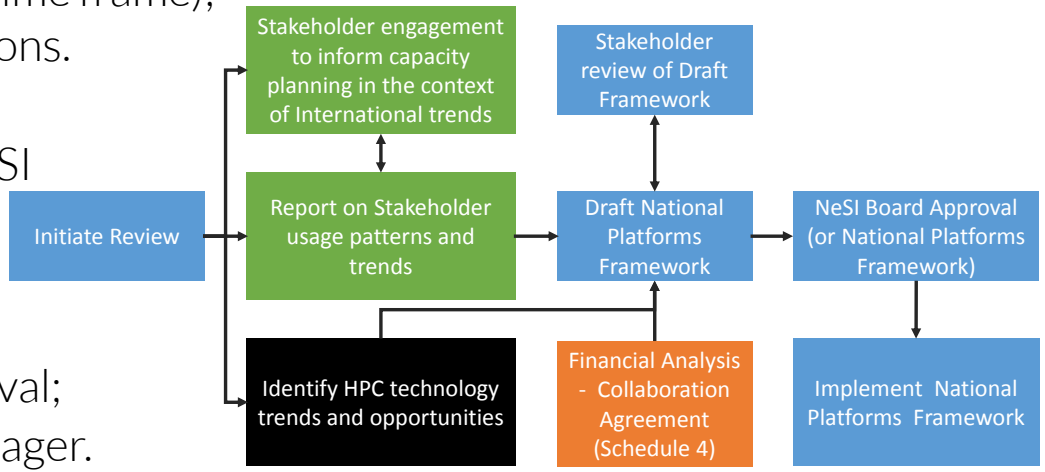

The 2015 Review

- NeSI HPC Platforms are approaching end of life:
  - Foster (UC BlueGene (IBM Power, Linux)), and P7 (IBM Power, Linux & AIX) commissioned in 2011;
  - Pan (UoA iDataPlex (Intel, Linux)), commissioned in 2011;
  - FitzRoy (NIWA P575/P6 (IBM Power, AIX)), commissioned in 2010.
- 2014 National Platform Framework proposes a two HPC Platform environment in the future:
  - Capacity platform at UoA, with Cloud bursting capacity on demand;
  - Capability platform at NIWA.
  - A focus on optimisation of fit-for-purpose use of the Platforms:
- NeSI Annual Plan.

Definitions: Capacity & Capability

- Capacity (e.g. Pan, P7) Application Domain:
  - 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);

- Capability (e.g. FitzRoy, Foster) Application Domain:
  - Large, highly coupled problems, which have high inter-processor / low latency communications requirements, and typically, very high I/O demands;
  - Tightly coupled problems that exhibit poor scaling properties require high performance processors;

HPC Compute 2013 – 15: By Allocation Class

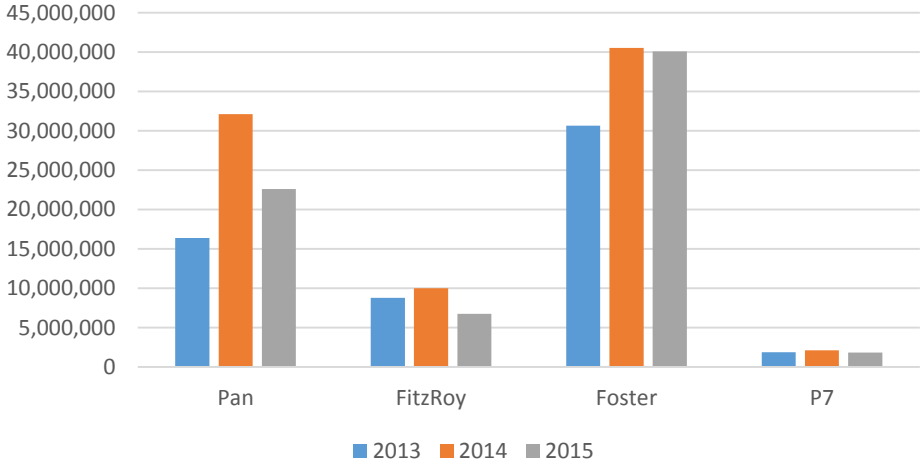

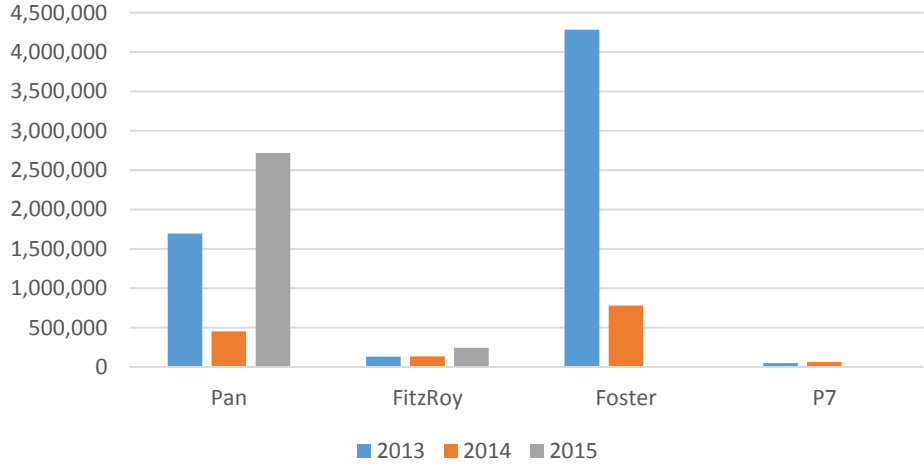

Total Compute 2013 – to date = **224,202,521** Core-h

Total Projects 2013 – to date = **911**

Target **60%** Institutional (UoA, Landcare, UoO, NIWA, UC)

40% for Merit, Post Grad, Proposal Development, Subscription

HPC Compute 2013 – 15: Project Metrics

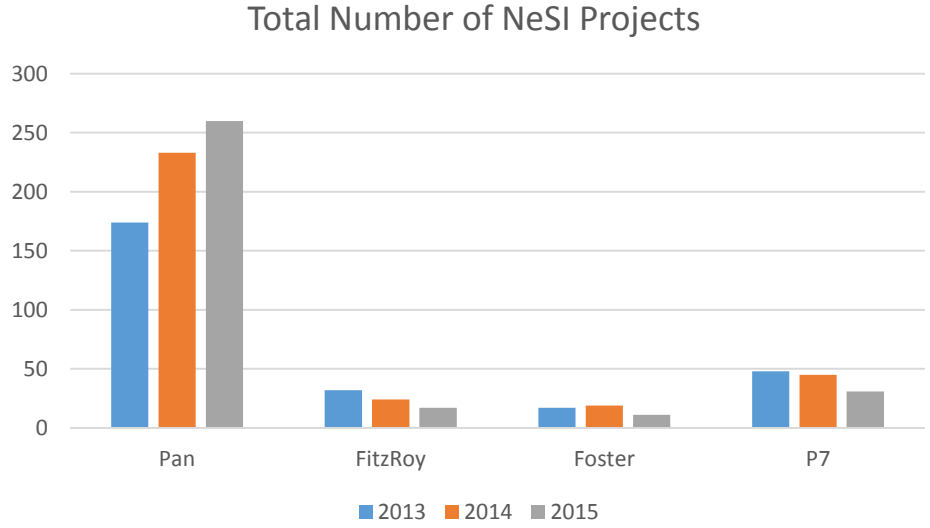

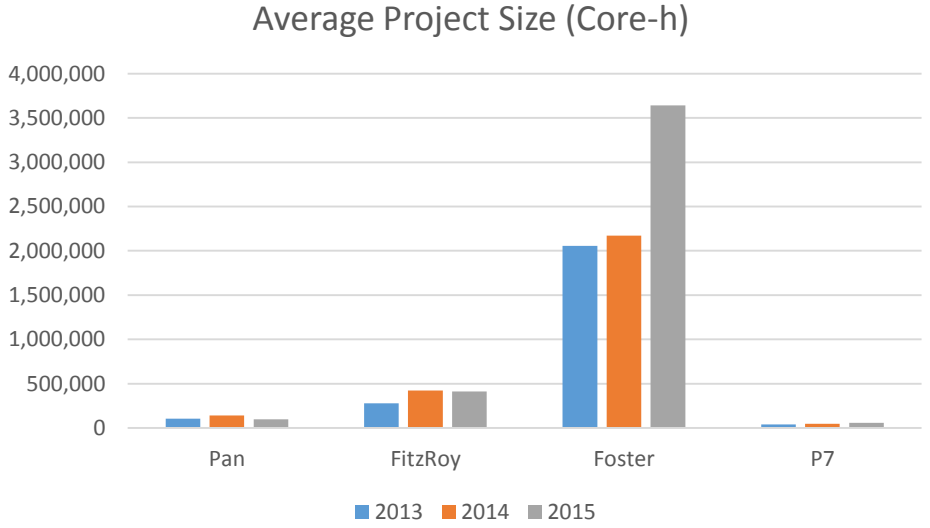

Total Compute 2013 – to date = **224,202,521** Core-h

Total Projects 2013 – to date = **911**

Looking to the Future: Research Needs for HPC

- The "Research Needs Survey" was used to gather information on:
  - Which research groups are dependent on HPC Services and what type;
  - International collaborations and dependencies (NSSI);
  - The HPC services that will be needed over a 3 – 5 year time frame to remain internationally competitive:
    - » CPU cores;
    - » Accelerators (GPGPU, MIC, FPGA);
    - » Data storage;
    - » Data Analytics (including HADOOP, SPARK, …?).
  - The types of Application that they are/will use and the scale;
  - Software codes that could be used in RFP benchmarking;
  - Gaps in NeSI's current service offerings.

Research Needs Survey: Overview

Responses:

- Auckland (18)
- Canterbury (2)
- Otago (5)
- Massey (1)
- Landcare (1)
- NIWA (18)
- NZGL (1)

Note: Not all respondents are using current platforms.

	Primary Research Domain	Responses	Platform		
			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**	**3 + (1)**

Headline Responses

- Data Services:
  - Faster methods to transfer large datasets between research groups;
  - Access to, and management of large datasets (e.g. to host reference datasets).
- HPC Compute and Analytics:
  - The big Earth Sciences and Astronomy researchers have a clear view of future needs, e.g.
    - Need for high performance cores and interconnects – for tightly coupled codes (not much use yet for GPGPUs or MIC architectures);
    - Very large Core-hour requirements for some planned research projects (O(100M Core-hours) per annum)
    - Large data output and storage (O (1PB) per simulation and the need for multiple simulations.

Headline Responses: Continued

- Researchers in Biomedical Sciences will also need access to large Capability Platform resources
- 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
- Use of MIC architectures (i.e. Knights Landing) in science codes that deliver performance improvements in time to solution is less clear.
- There is a substantial need for Capacity services – the long tail of HPC.

- Data Analytics
  - 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
  - In part these will be driven by the need to analyse PB scale datasets

Headline Responses: Continued

- Visualisation
  - Little comment – but server/GPU based visualisation will be important (i.e. don't move data!).
- Platforms Operations:
  - Easier to transition between the two platforms (user environment / data);
  - Better transparency / management of job queues;
  - Run on "fit-for-purpose" Platforms;
  - Data management.
- Need good planning around decommissioning activities.

Platform Replacement Design Considerations

- Make it easy for users to develop and run research workloads/jobs and apply data analytics tools on either/both platforms (minimise diversity);
- Fit-for-Purpose platforms that meet researcher needs;
- High level of interoperability/commonality of management and monitoring systems;
- Leverage step changes in technology: processors, software environment, storage;
- Transparent management of data on tiers (from Flash ↔ disk ↔ tape);
- Fastest time to solution;
- High reliability and availability;
- Minimise Total Cost of Ownership;
- Access to standard "big data" tools (e.g. Hadoop, SPARK)?

Draft National Platforms Framework (2015)

- 2016:
  - Decide on the role of NeSI Cloud-Burst services in the context of the Framework;
  - Agree Data Services strategy and feed into Platform replacement design;
  - Agree and design Data Analytics capabilities / services, and feed into Platform replacement design;
  - Develop and issue RFPs for both Capacity and Capability Platforms (i.e. Design and specify solutions that will meet NeSI Goals and Objectives, informed by Researcher Input);
  - Analyse RFP responses and select vendor(s).
- 2017:
  - Contracting, installation, acceptance testing, configuration, commissioning, transition to operations, decommission old platforms;
  - Optimize services .
- 2018:
  - Optimize services.

Extra Slides

NeSI.2 Access Policy: Summary

Merit (No cost):

- Highest priority allocations, awards in this class are given priority access to Consultancy
- Projects are required to provide evidence of an existing peer reviewed national or institutional award
- Project allocated to best fit-for-purpose Platform
- Grants up to one year, renewable.

Institution (Cost):

- Access a "block" allocation (Core-h and Consultancy);
- Service Governance by Institution, Technical Assessment by NeSI.

Proposal Development (No cost);

- Fast access, to learn about HPC etc..,
- Small allocation (1000 Core-h), time bound (1 month).

Post Graduate (No cost):

- Available to post-graduate students – working on an approved research programme);
- Lowest priority;
- Grants up to one year, renewable.

NOTE: "Merit" Access includes Merit (as defined above), Proposal Development, Post Grad and Subscription usage