



The NeSI HPC Compute and Data Analytics Service: New Systems, Capabilities & Science

Michael Uddstrom (Platforms Manager)

Blair Bethwaite (Solutions Manager)

New Zealand eScience Infrastructure

Design Choices (NIWA + NeSI Co-investment)

- National data-centric research and operational computing environment:


- Single site: NIWA Greta Point, Wellington;
- HPC Cluster: Low core count jobs
- Supercomputer: Large core count jobs & NIWA forecasting;
- GPGPUs;
- DTN service;
- Shared (large) high performance filesystems;
- Secure (MFA).

- New user services:

- Virtual Labs (interactive access to data);
- Remote visualisation;
- Private cloud;
- Advanced data analytics;
- Librarian (move data to, and from offline).

- Disaster Recovery:

- Copy of NeSI research data and backups;
- NIWA Supercomputer.



Reprising last year's: Next Steps

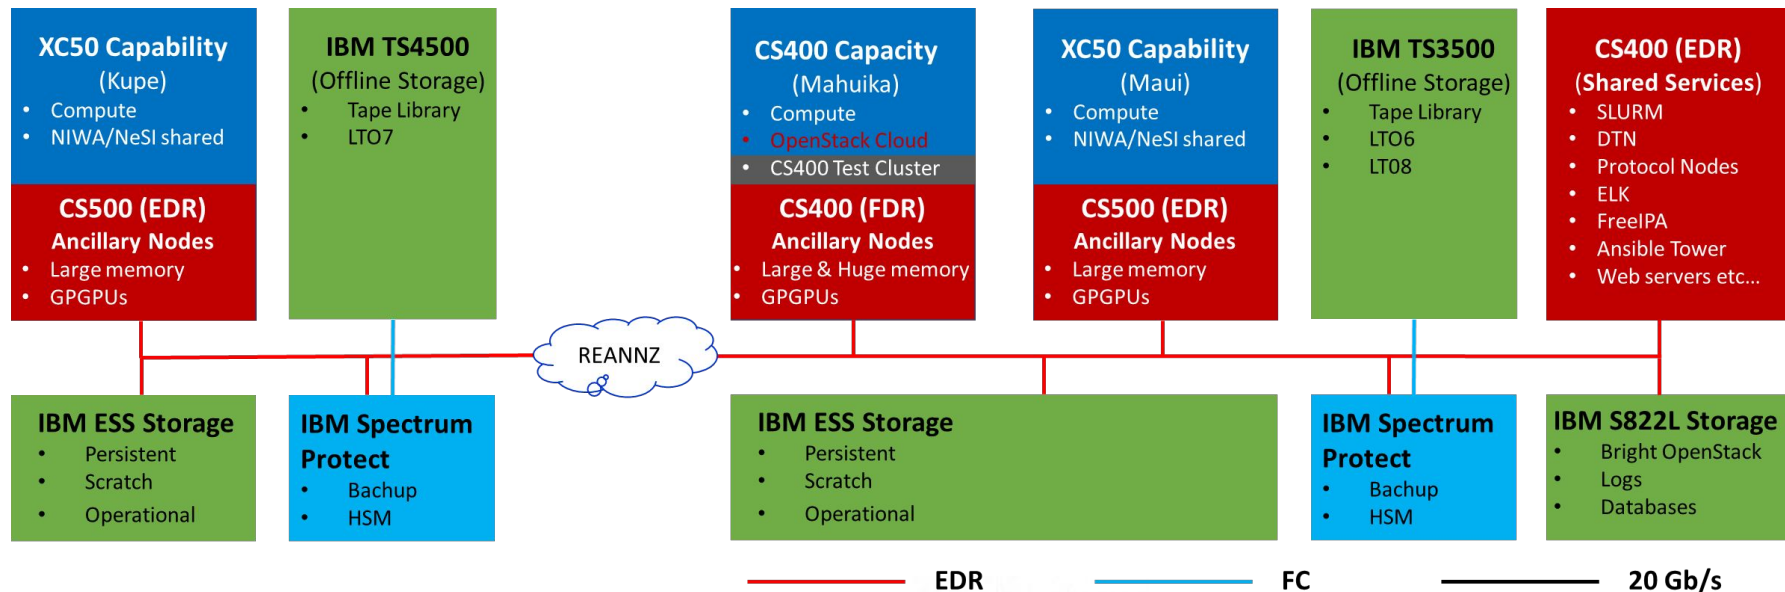
“Predictions”

- ~~Greta Point plant upgrade~~ (Housing, Power, Cooling) 28-Feb-2018 ✓
- Begin install of Mahuika and Maui 26-Feb-2018 ✓
- Acceptance Testing, Configuration and Tailoring 31-Mar-2018 ✓
- Target to Transition all users (Pan, Kupe) to HPCF 01-May-2018 ✗
- Implement HPC Compute and Analytics Service design 01-May-2018 ✓
 - Fit for purpose website content:
 - » Providing prospective users with relevant information. ✓
 - » <https://support.nesi.org.nz>
 - Improved project application, review and allocation management; ✓
 - Deeper user support; ✓
 - New services:
 - » Virtual Workstations/Labs, Remote Visualisation, Hybrid Cloud, HSM, etc. ✓ / ✗

High Performance Computing Facility

NIWA/NeSI – Tamaki Data Centre

NIWA/NeSI – Greta Point



CRAY

IBM

Bright Computing

High Performance Computing Facility (1.7PF)

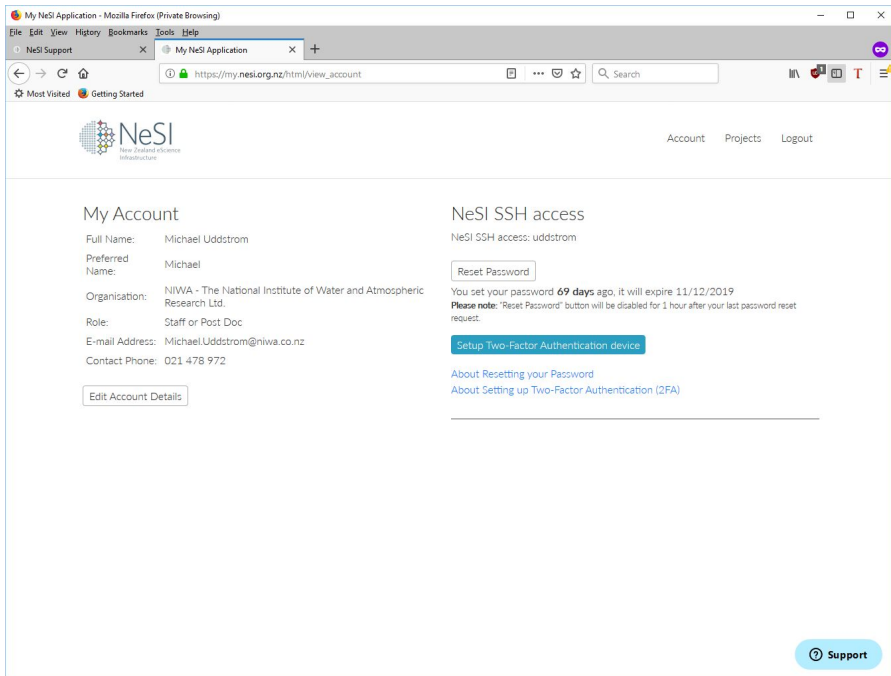


Mahuika (9,064 BRW cores)
Maui_Ancil (1,120 SKL cores)

IBM ESS Shared Storage
10.1PB (~150GB/s)

Maui (18,560 SKL cores)
(1.4 PF)

New NeSI Access & Support Pages



My NeSI Application - Mozilla Firefox (Private Browsing)

https://my.nesi.org.nz/html/view_account

NeSI

Account Projects Logout

My Account

Full Name: Michael Uddstrom
Preferred Name: Michael
Organisation: NIWA - The National Institute of Water and Atmospheric Research Ltd.
Role: Staff or Post Doc
E-mail Address: Michael.Uddstrom@niwa.co.nz
Contact Phone: 021 478 972

Edit Account Details

NeSI SSH access

NeSI SSH access: uddstrom

Reset Password

You set your password **69 days** ago, it will expire 11/12/2019

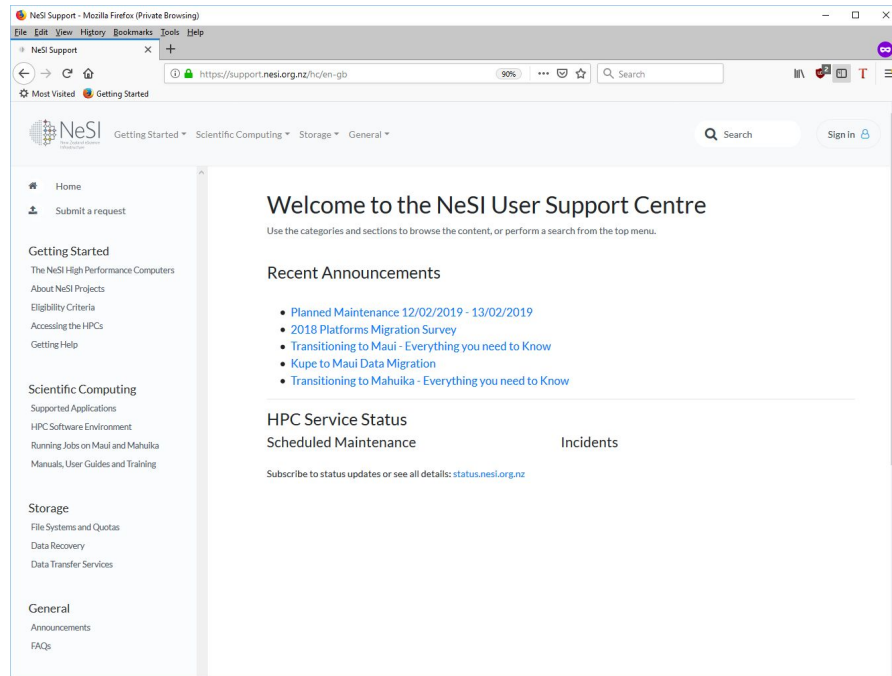
Please note: "Reset Password" button will be disabled for 1 hour after your last password reset request.

Setup Two-Factor Authentication device

About Resetting your Password
About Setting up Two-Factor Authentication (2FA)

Support

<https://my.nesi.org.nz>



NeSI Support - Mozilla Firefox (Private Browsing)

https://support.nesi.org.nz/hc/en-gb

NeSI

Getting Started Scientific Computing Storage General

Search Sign In

Welcome to the NeSI User Support Centre

Use the categories and sections to browse the content, or perform a search from the top menu.

Recent Announcements

- Planned Maintenance 12/02/2019 - 13/02/2019
- 2018 Platforms Migration Survey
- Transitioning to Maui - Everything you need to Know
- Kupe to Maui Data Migration
- Transitioning to Mahuika - Everything you need to Know

HPC Service Status

Scheduled Maintenance Incidents

Subscribe to status updates or see all details: status.nesi.org.nz

Getting Started

The NeSI High Performance Computers

About NeSI Projects

Eligibility Criteria

Accessing the HPCs

Getting Help

Scientific Computing

Supported Applications

HPC Software Environment

Running Jobs on Maui and Mahuika

Manuals, User Guides and Training

Storage

File Systems and Quotas

Data Recovery

Data Transfer Services

General

Announcements

FAQs

<https://support.nesi.org.nz>

NeSI HPC (X86) Compute Resources (per


Allocation Class	Mahuika (excluding cloud) (core-hs)	Mahuika (cloud) (core-hs)	Mahuika (Ancillary Nodes) (core-hs)	Maui (node-hs)	Maui (Ancillary nodes) (node-hs)
Merit	12,580,834	504,922	1,122,048	362,870	1,122,048
Proposal Development					
Post Graduate					
Subscription	12,580,834	504,922	1,122,048	362,870	1,122,048
Collaborator	37,742,399	1,514,764	3,366,144	1,088,611	3,366,144
Total (core-hs)	62,904,067	2,524,608	5,610,240	75,574,040	5,610,240

NeSI HPC (GPU) Compute Resources (per

Allocation Class	Mahuika (GPU-hs)	Maui (GPU-hs)	Comment
Merit	13,736	8,585	1 × GPU-h is equivalent to 3,584 Cuda Core-hours.
Proposal Development			
Post Graduate			
Subscription	13,736	8,585	
Collaborator	41,206	25,754	
Total (GPU-hs)	68,678	42,924	Equivalent to 400 million Cuda Core-hs per annum

- Resources: <https://support.nesi.org.nz/hc/en-gb/articles/360000201756-Total-HPC-Resources-Available>
- Pan & Mahuika compared:
<https://support.nesi.org.nz/hc/en-gb/articles/360000204715-Pan-Mahuika-Compared>
- FitzRoy & Maui compared:
<https://support.nesi.org.nz/hc/en-gb/articles/360000204855-FitzRoy-M%C4%81ui-Compared->

System Activity: Mahuika job dimensions

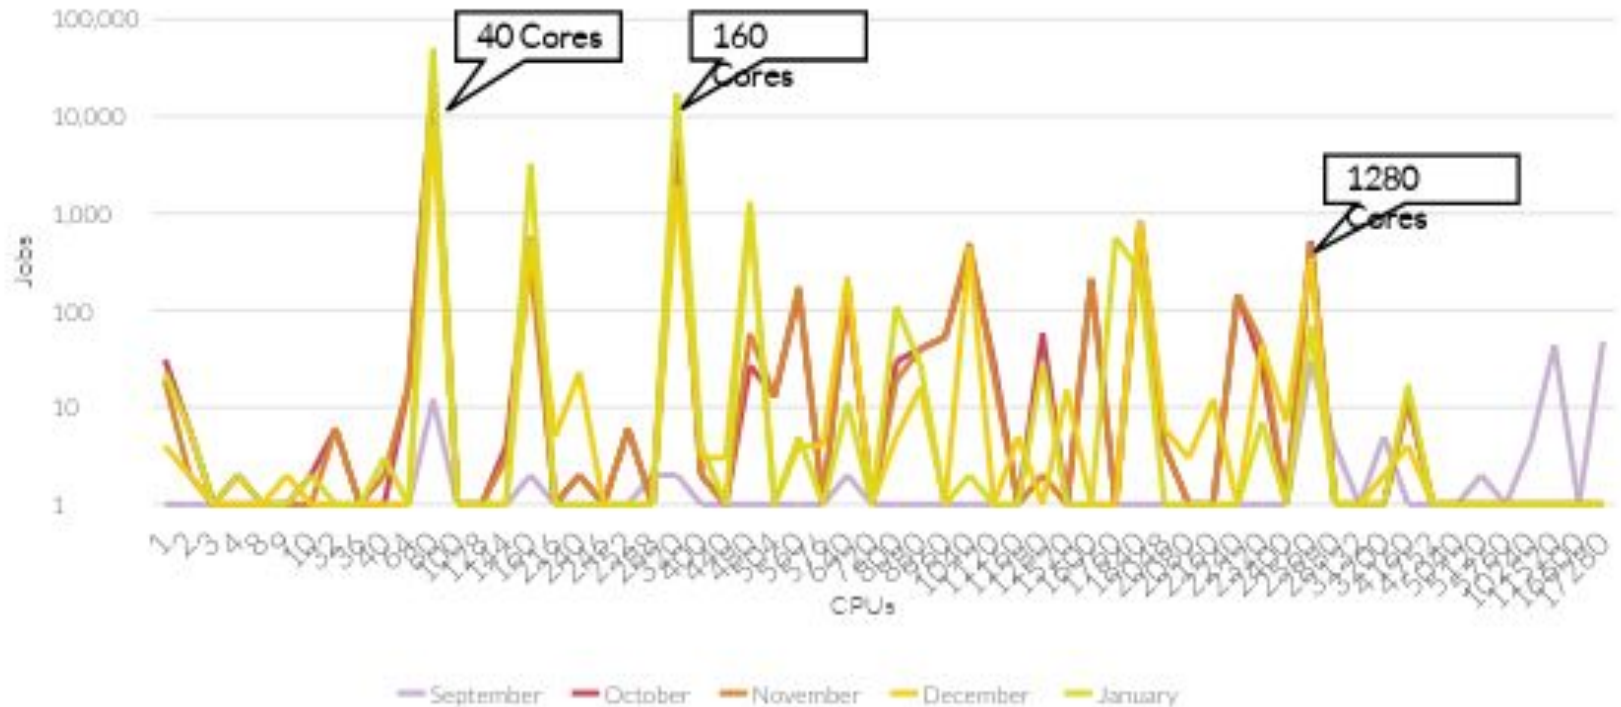


System Activity: Mahuika

- Some points to note:
 - On Mahuika all but the huge memory node have 2×18 core sockets (i.e. 36 cores);
 - Hyperthreading is enabled \Rightarrow Slurm can schedule **72 Slurm CPUs** per node;
 - On Mahuika base compute nodes have **~100GB** of memory available;
 - » Users typically request 50% more than they use
 - On Maui the equivalent values are 20, 40 and 80.
- On Pan hyperthreading was disabled and nodes had 12, 16, (20) and 24 cores;
- It looks like job scripts **may not have been updated to reflect the new processors.**

System Activity: Maui


Number of Jobs and Slurm CPUs allocated per month - Maui



Why a Supercomputer to Forecast the Weather?

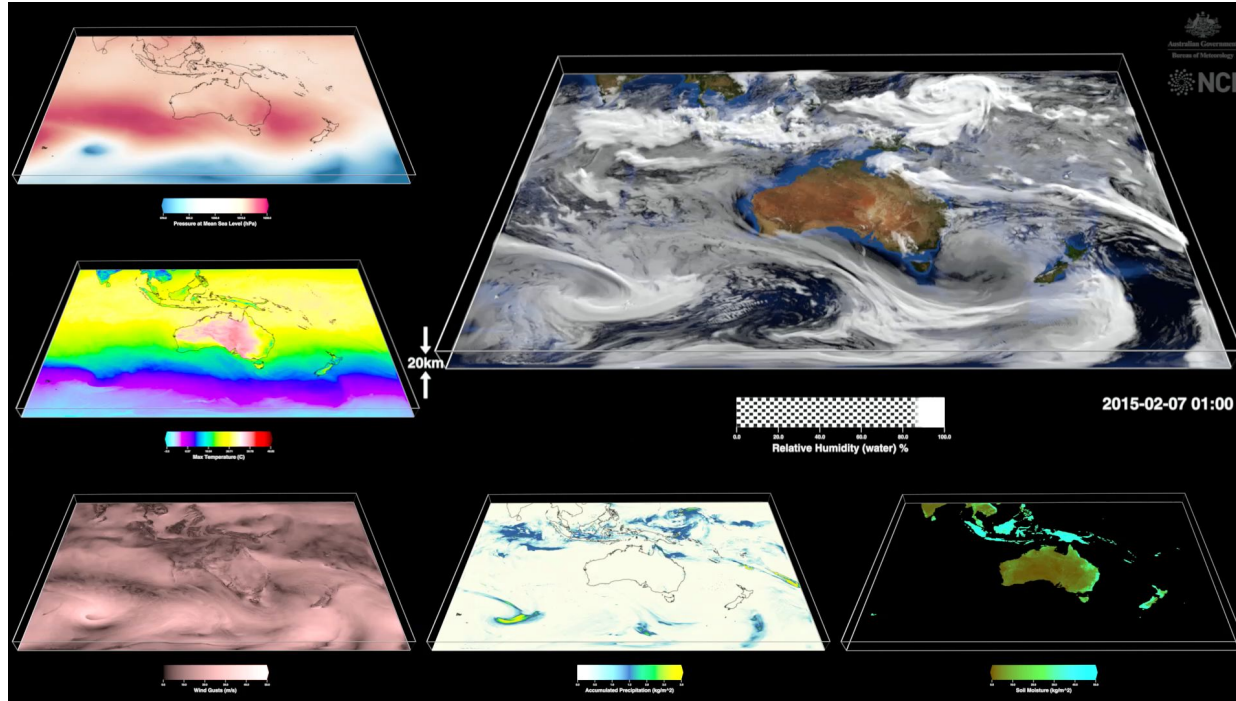
The goal of numerical weather prediction (NWP) is to predict the future state of the atmosphere given:

- Observations of its current state;
- An atmospheric simulation model: a set of numerical approximations to a closed set of equations that govern atmospheric motion
- The NZCSM NWP model (1.5 km grid, convection resolving)
 - 567,000,000 locations/variables at every timestep;
 - Dynamics timestep = 60 sec:
⇒ 1.43×10^{12} calculations for a 42 h forecast (dynamics only);
- Cray XC50 1240 cores, 110 min



Preparing for the future... First Look Back

More than
500TB of data
(and counting)



Credit: Drew Whitehouse (NCI) – BARRA-R simulation of cyclone Yasi in Jan 2011

Virtual Laboratories – what are they?

- New feature of the HPC and Data Analytics Service;
- Project-specific dedicated and persistent interactive environments;
- Native (high-performance) access to HPC filesystems;
 - Single stream access: 5(W) – 6(R) GB/s (Aggregate Bandwidth 150 GB/s)
- Tightly integrated:
 - Identity and Access Management;
 - Software;
 - Filesystems;
 - HPC job submission and management;
 - But no superuser.
- Delivered via OpenStack with scalable capacity that can grow/shrink over cluster *ancillary* nodes (i.e., everything not a standard compute node)

Virtual Laboratories – why?

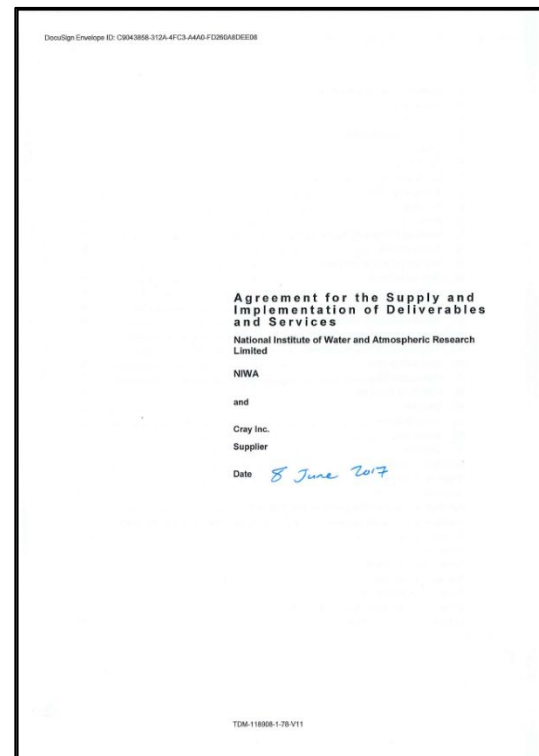
- Exploratory Data Analysis, and development:
 - Remove resource hungry workloads away from shared login nodes;
 - (May be) Scientific domain specific.
- HPC facing persistent services e.g. databases;
- HPC workflows triggered by external data ingest and event processing;
- Computational steering and workflows;
- Remote visualisation.

Virtual Laboratories – I want one!

- We are still developing and refining the delivery model...
- For now please contact NeSI support (in future they will be available through my.nesi);
- No self-service option yet - possible future addition depending on use-cases and interest;
- Initial NeSI pilot VLs underway working on NWA, Genomics Aotearoa and Manaaki Whenua Landcare Research use-cases
- Looking for more willing beta users!

Reminder: “Good Things” take time....

- 2016
 - National Frameworks Review approved: 03/16
 - Co-Investment: NIWA \$15.5M; UoA+UoO+MWLR \$4.8M
 - Collaborative (NeSI/NIWA design process)
 - One RFP for 3 HPCs released: 12/16
- 2017
 - Cray awarded contract: 07/17
 - NIWA's Kupe system operational: 12/17
- 2018
 - Pan and FitzRoy replacements accepted: 03/18
 - Operational service available: 08/18 & 09/18
- 2 years 6 months (1 year from Contracting to Operations)



Summary

- The NZ research community has access to leading edge HPC and Supercomputer resources.
 - Providing access to services equivalent or better than any of our international collaborators
- The Platforms and HPC Compute and Data Analytics Service refresh has delivered
 - >10× more compute performance
 - 20 to 150× more filesystem performance
- New services
- Already – we are seeing new science results (e.g. QuakeCore, NZESM)
- Via NIWA – the ability to deliver resilient 24 x 7 operational services

NeSI @ eResearch NZ - Talks & Workshops:



~~Monday 18 Feb~~

~~2:10 - 2:30 pm - Understanding research drivers for NZ's advanced research computing~~

~~2:30 - 2:50 pm - How NeSI helps Manaaki Whenua Landcare Research monitor land cover changes~~

~~3:30 - 3:50 pm - NeSI Futures~~

~~4:30 - 5:30 pm - Training Community BoF~~

~~4:50 - 5:10 pm - Catering to domain (Genomics) specific eResearch needs~~

Tuesday 19 Feb

11:00 - 11:20 am - The NeSI HPC Compute and Data Analytics Service

11:00 am - 12:30 pm - Open Space Session - BYO topics!

11:20 am - 11:40 am - Deploying a Globus endpoint in an NZ institution

1:30 - 1:50 pm - Visualization capabilities of NeSI's new high performance computers

1:30 - 1:50 pm - A day in the life of NeSI's Apps Support

1:50 - 2:10 pm - NeSI and your data: Scalable storage

1:50 - 2:10 pm - Research Software Engineering (RSE): What's in a name?

Tuesday 19 Feb (cont.)

2:10 - 2:30 pm - Kicking On: Scaling new data services at NeSI

2:30 - 2:50 pm - Insight into the new NeSI platforms

3:30 - 4:30 pm - (Inter)national collaborative research infrastructure strategies BoF

3:30 - 4:30 pm - Research Software Engineering BoF

4:30 - 5:30 pm - Research Cloud NZ BoF

Wednesday 20 Feb

11:10 am - 4:00 pm - Hacky Hour / Bring Your Own Code Workshop