



nectar

National eResearch Collaboration Tools and Resources

NeCTAR is an Australian Government project conducted as part of the Super Science initiative and financed by the Education Investment Fund. The University of Melbourne has been appointed the lead agent by the Commonwealth of Australia, Department of Innovation, Industry, Science and Research.

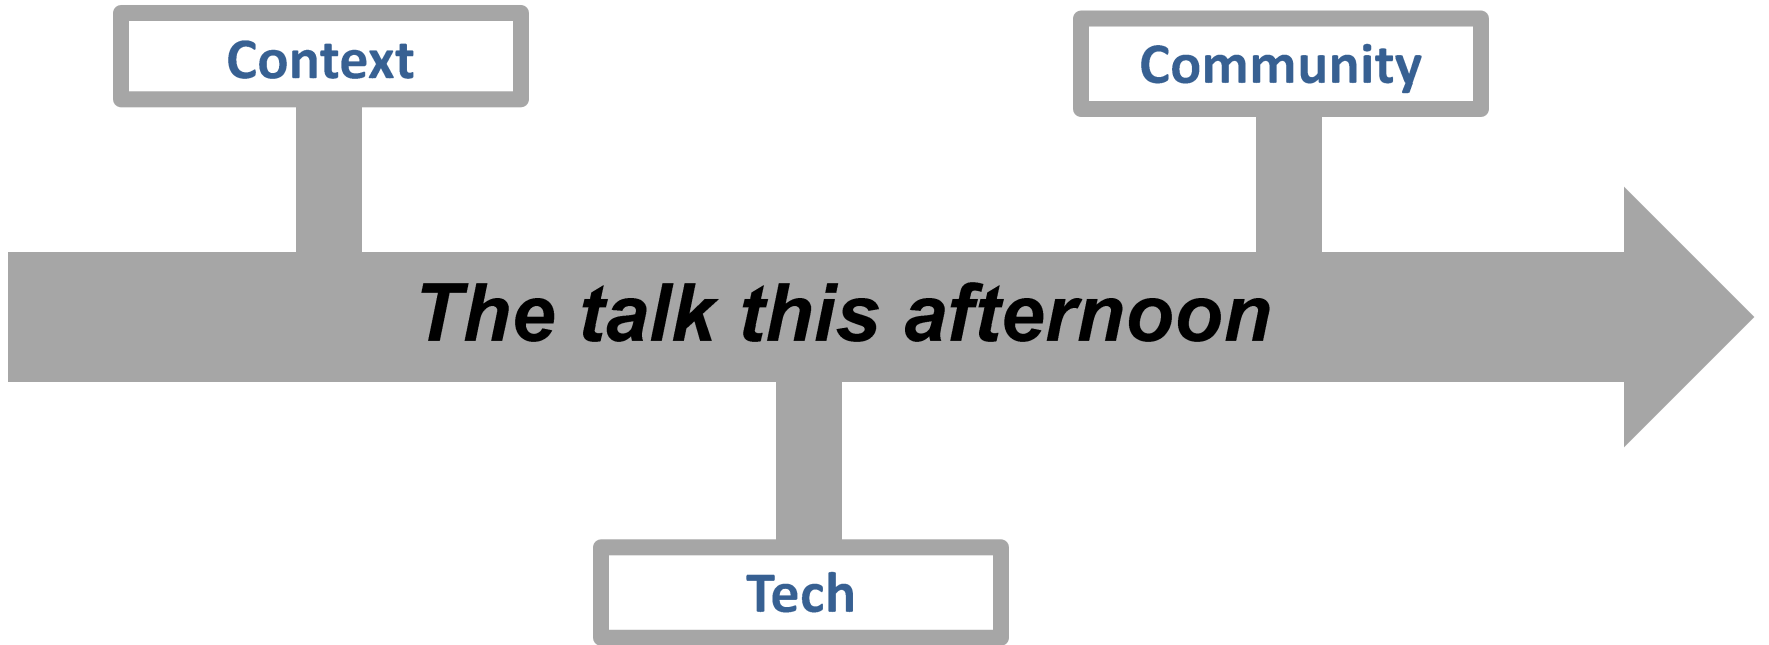
Objectives: to enhance research collaboration through the development of eResearch infrastructure.

Context

Community

The talk this afternoon

Tech



eResearch

The process of research can readily be described in term of how data flows

Compute

Computational modeling, data analysis, etc. (APAC, NCI, Pawsey)

Combine

Use tools & apps to work remotely and collaboratively - NeCTAR

Keep

Keep data and observations, describe, collect, share, find and re-use them – RDSI, ANDS

Access

Overall Timeline - Infrastructure Extension

Development of the Australian Access Federation *AAF*

Previous Development of the AREN | AREN Extensions *NRN*

Previous Peak Computing | Peak Computing *NCI* | NCI Peta scale

New Peak Computing | Pawsey Peta scale

Research Tools *ANDS, ARCS*

Research Tools, Workflows & Cloud Services *NeCTAR*

Collaboration, data, grid *ARCS*

Research Data Commons *ANDS*

Data Storage Services *RDSI*

NCRIS Announced

Super Science Announced

Road map

Road map

2005 | 2006 | 2007 | 2008 | 2009 | 2010 | 2011 | 2012 | 2013 | 2014

PfC Consultation

HPC \$26M NCI
Collab \$22M ARCS

Data \$24M ANDS

Data \$47M ANDS

HPC \$80M Pawsey

HPC \$50M NCI

Cloud \$47M NeCTAR

Networks \$37M NRN

Storage \$50M RDSI

1 AUD
~1 USD
~6.7 CNY

NeCTAR has four program areas and two operational national services

Research
software



Virtual laboratories



eResearch Tools

Computational
platforms



Research Cloud



National Server
Program

This cloud... why build it ourselves?

- 1. Proximity – the honeypot** – infrastructure attracts community
- 2. Local infrastructure** is more responsive to research needs
- 3. Service offering and usage modes** suitable for research
- 4. Locality** to instruments, research networks and other infrastructure.
- 5. Data sovereignty**


NeCTAR Research Cloud

An OpenStack Cloud

- Any Researcher, Any Discipline, Anywhere

A single national cloud

- Up to seven 'nodes'
- \$1.5M per node
- ~4000 cores/node

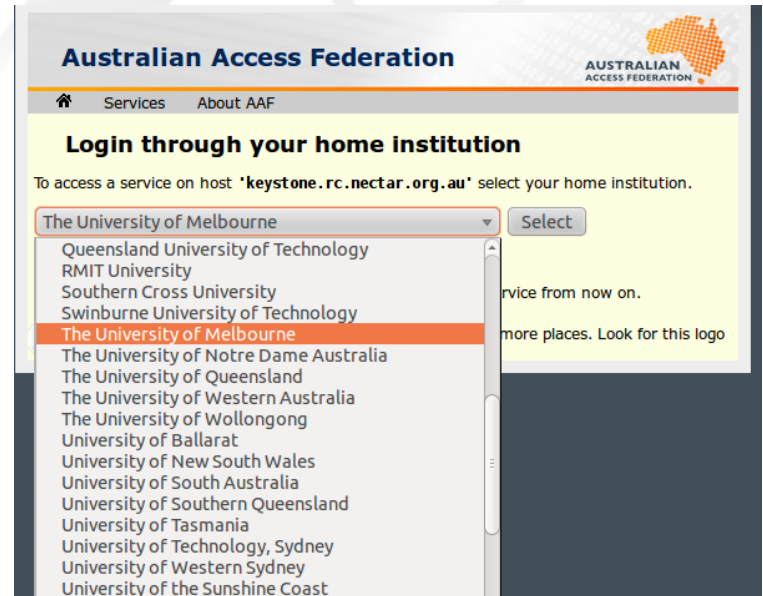


Supporting Inward and Outward Federation

- EC2/S3 API, OpenStack API

Low barriers to entry

- Single Sign On
- Two cores free for 3 months
- Upload own VM Images
- Responsive Support
- Cater to different levels



The screenshot shows the Australian Access Federation (AAF) website. The header includes the AAF logo and navigation links for 'Services' and 'About AAF'. The main content area is titled 'Login through your home institution' and instructs users to select their home institution to access services on the host 'keystone.rc.nectar.org.au'. A dropdown menu is open, listing various Australian universities, with 'The University of Melbourne' selected. A 'Select' button is visible next to the dropdown. The page also contains text about service availability and a note about logos.

Cloud platform for innovation & open research

A platform for hosting, deploying and sharing research software infrastructure

Supporting collaboration and innovation in research software and services

Early, rapid deployment and sharing of research applications on a national scale

Reducing barriers to success *and reducing the cost of failure!*

Research computation which complements HPC investments

Complements other initiatives as well, research data management, petascale storage, networks, etc.

Build to a research spec: researchers work 24/7, globally, collaboratively across boundaries

Project is funding applications & virtual labs to sit atop the cloud

Community & Innovation

The research cloud as a platform for innovation: Give researchers time to do exactly that

Conversation shifts: sharing code to sharing data and sharing servers openly

Communities (discipline-based, tech-based, etc.) will form around the cloud, as we work out how to use it more effectively

The conversations we're having in 2012 (cloud) will be *very* different in 2013 (research)

Real research

Collaboration tools e.g. Sakai, Wordpress and MediaWiki



Particle Physics simulations powered by a distributed computing workflow tool to span Grid, Cloud and high-performance computing

Large-scale genome analysis using the Galaxy interactive platform

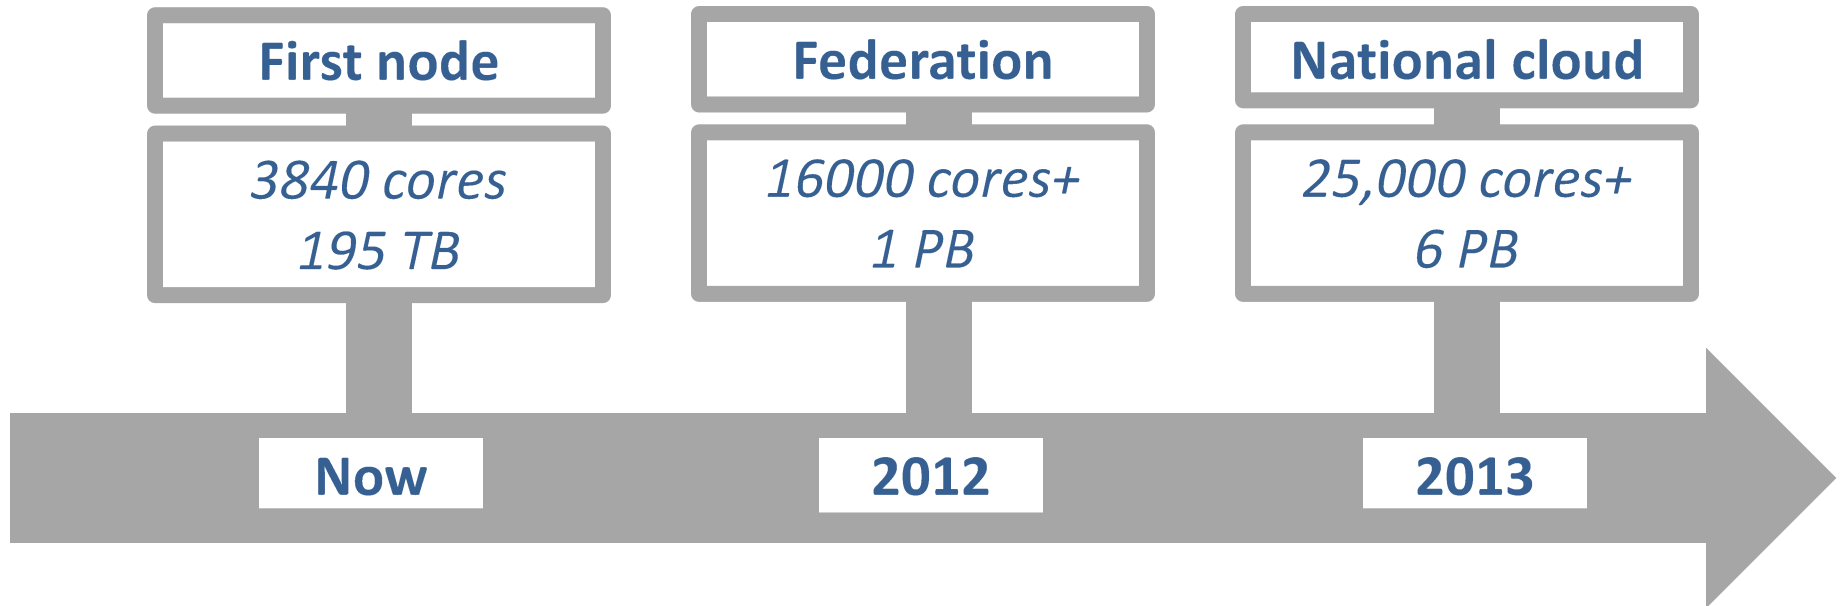
Digital Humanities – Humanities Networked Infrastructure (HuNI)

Disaster Management - BushfireConnect.org

Urban Research and geospatial information platforms



The story so far



Specifications – University of Melbourne node

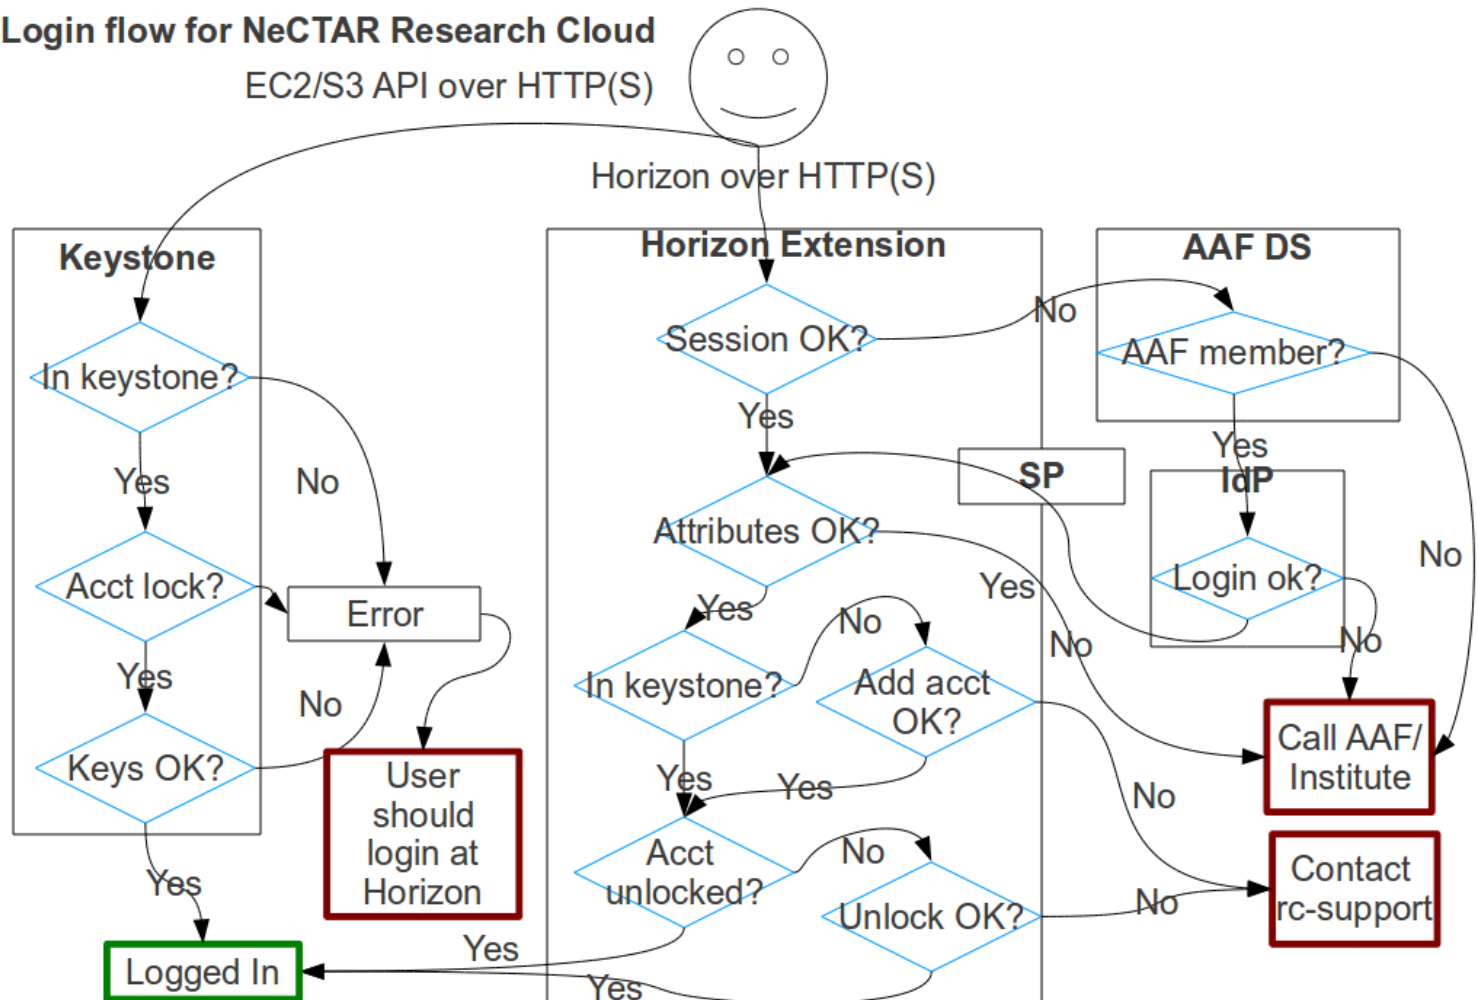
- OpenStack Essex/Stable (+~5% Folsom backport)
 - Ubuntu 12.04 LTS, Puppet
- Hardware
 - 336 cores – 48 Core Dell R815s
 - 3840 cores – 160* 24 core, 128GB, 10Gbit/s Xenon Quad2U
 - 195TB – HP DL180G w/ DL2000 @ 24TB/node
 - 146TB – Dell R715 w/ MD1200 @ 24TB/node
 - 10Gbit/s – Cisco Nexus (2232, 5596, ...)
 - Hitachi HNAS/BluARC – 100TB for running VMs

Working with OpenStack

- Shibboleth login to dashboard

Login flow for NeCTAR Research Cloud

EC2/S3 API over HTTP(S)



Working with OpenStack

- Dashboard Allocation Form



USER DASHBOARD

Manage Compute

Overview

Credentials

Instances

Images

Snapshots

Keypairs

Security Groups

Allocation Request

Allocation Request

Description

This form allows you to request a project specific allocation on the research cloud.

Note: This application form is currently for use by preferred Virtual Laboratories, Early Approval Virtual Laboratories and Research Tools projects. This form will be made m

Required fields are marked with a *.

Once you have submitted the form your request will be begin its journey through the allocation process. You will receive e-mails at the address specified by your [AAF](#) [IdP](#)

* Project name



A human friendly descriptive name for your project.

Contact e-mail

fifieldt@unimelb.edu.au



The e-mail address provided by your IdP which will be used to communicate with you about this allocation request.

Working with OpenStack

- Network counting
 - Our ISP AARNet has two traffic tariffs:
 - On-net (universities, research institutes anywhere in the world) – which costs **\$0**
 - Off-net – which costs **\$\$**
 - Each AARNet CPE can provide a netflow feed
 - We use this, hooked up to OpenStack to associate netflows to our users
 - Policy: 1GB/core/month of Off-net, then <action>

Working with OpenStack

OPENSTACK
DESIGN SUMMIT & CONFERENCE

APRIL 16-20



Compute Cells

Chris Behrens

cbehrens@codestud.com

comstud@IRC

<http://comstud.com/FolsomCells.pdf>


[http://etherpad.openstack.org/
FolsomComputeCells](http://etherpad.openstack.org/FolsomComputeCells)

Working with OpenStack

Minimal services

API cell	Child compute cell
AMQP broker (for RPC)	AMQP broker (for RPC)
DB	DB
nova-cells	nova-cells
nova-api	nova-scheduler
	nova-network
	nova-compute

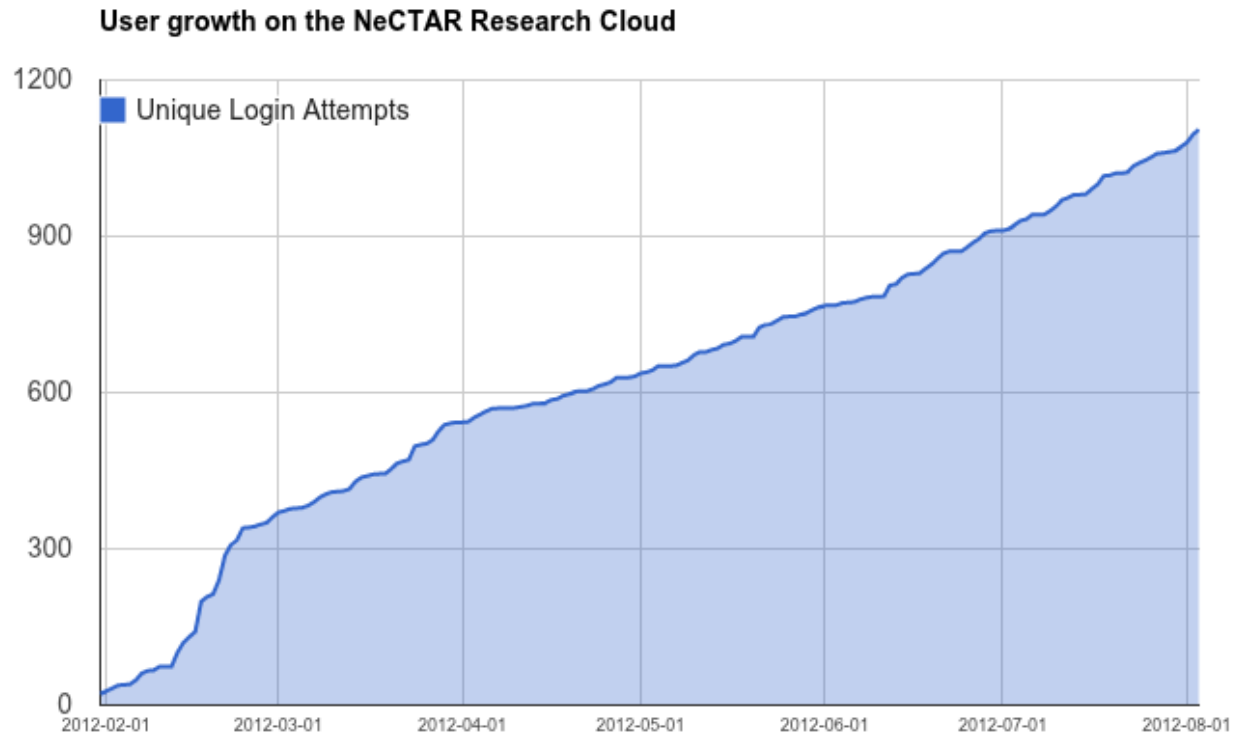
Working with OpenStack



Build community at all levels

Research Cloud Users

- #NADojo full day dev workshops
- Starting monthly meet-ups



Data & Cloud DevDays

#NADojo : Melbourne *“I used rely on mainframes to run my experiments. One time I worked Saturday night because it was the only time I could get access. I can’t build a computer, but the fact that there’s instant access alone is enough for me to use the cloud”*



#NADojo Brisbane *“These types of forums are import not only to bring us together and share out experiences. These are new technologies, and things like this give us a chance to get a step up, while meeting the community.”*



Build community at all levels

Australian OpenStack Community

- Sponsored by Aptira – thanks for the beer!
- Australia is large, but sparsely populated
 - East-West is ~4 hours **flying** time
 - ~Size of US, with population 25 million (3 people/km²)
- One community, meet-ups in multiple cities
- Meetings so far in Melbourne(x2), Sydney (x3), Brisbane, attended by 35-60 people each.
- Next meetings: Adelaide, Hobart, Perth



APTIRA AND RACKSPACE – NEW FRIENDS BY OPENSTACK



THE PEOPLE OF AOSUG



SYDNEY AND MELBOURNE INAUGURAL MEETUPS DEC 13 & JAN 17



THE SYDNEY OZSTACKERS MEET JAMES WILLIAMS



MARCH 26 OPENSTACK TECHNICAL PRESENTATION MEETUP



JOHN DICKINSON – OUR SWIFT BIG BROTHER?



1. Timing of sessions is important.
2. Reach out wherever you can.
3. Get vendors involved.
4. Check the venue, and check it again.
5. Under estimate RSVP numbers.
6. Engage the serial RSVPers, (thanks Martin!)
7. Beware of MeetUp.com calendar appointments!
8. Sponsors and contributors should have priority.



nectar

National eResearch Collaboration Tools and Resources

Questions?

NeCTAR is an Australian Government project conducted as part of the Super Science initiative and financed by the Education Investment Fund. The University of Melbourne has been appointed the lead agent by the Commonwealth of Australia, Department of Innovation, Industry, Science and Research.

Objectives: to enhance research collaboration through the development of eResearch infrastructure.