

10 years and counting of the ARDC Nectar Research Cloud

19 October 2022
15.50-16.10

PRESENTED BY
Carmel Walsh
Paul Coddington

Australian Research Data Commons





Nectar Founder
Glenn Moloney
Nectar Director 2010 - 2018

NECTAR PROGRAM

- **National eResearch Collaboration Tools & Resources**
- Program started in 2010
- National Research Cloud
- Virtual Laboratories - ARDC Platforms
- Visionary approach to eResearch
- Cloud computing was a new technology
 - very little use in research
 - no commercial cloud sites in Australia

Origins of Nectar



FOUNDING PRINCIPLES OF NECTAR:

- ***Cloud based infrastructure to enable cross-institutional research collaboration for Australian researchers***
- ***Enable researchers to store, access, share and analyse data***

TECHNOLOGY - *Why OpenStack?*

- Want to provide cost-effective computing infrastructure and services **at scale**
- **Self-service** capability to enable users to quickly provision resources
- 2010 OpenStack created as a collaboration between NASA and Rackspace
- OpenStack was chosen in 2011 from a number of competing cloud products
- **Nectar was the first in the world** to adopt OpenStack for a distributed research cloud, helping to pioneer large-scale deployment of OpenStack, particularly for research
- **OpenStack was a good choice** - now the leading open source cloud software, used by thousands of organisations and many research clouds
- One of the **largest contributor bases** of any open source software
- Enables us to **adapt and grow** with changing researcher needs
- We benefit from, and contribute to, a **large international community**

OPENSTACK USERS

- 25M cores in production under OpenStack in 2021 (66% growth from 2020)
- Several organisations have > 1M cores
 - Walmart, Yahoo!, Volkswagen, Nike, Bloomberg, Target, Cathay Pacific
 - China UnionPay, Tencent, Postal Savings Bank of China - have ~500M users
 - AT&T, Verizon, Vodafone, China Mobile, Deutsche Telecom, Telstra
 - Several large commercial cloud providers - especially in EU and China
 - Australian commercial cloud providers - DUG, Firmus Supercloud, Vault
 - CERN, Harvard, MIT, Edinburgh, Oak Ridge National Laboratory

BUILDING A FEDERATION

Call for nodes

Two RFPs for nodes
in November 2011 and
May 2012.

Independent expert
panels to decide on
Nectar nodes and Virtual
Labs

Expectations

Work with the lead node
(UoM) / Core Services to
operate as a national
federation within an
evolving architectural
framework (OpenStack)
and standard operations,
maintenance and support
mechanisms

Criteria

Operational experience,
research community
engagement, support
models, training, etc.

Align with existing
eResearch infrastructure
including RDSI storage
nodes

BUILDING A COMMUNITY

- Training
- Tech/Ops Workshops
- Distributed Helpdesk
- Technical support to virtual labs & ARDC Platforms



We started with Dojo style technical training and train-the-trainer events around the country, introducing technical and user support staff to Nectar

- Sam Morrison, Technical Lead, Nectar Core Services since 2012



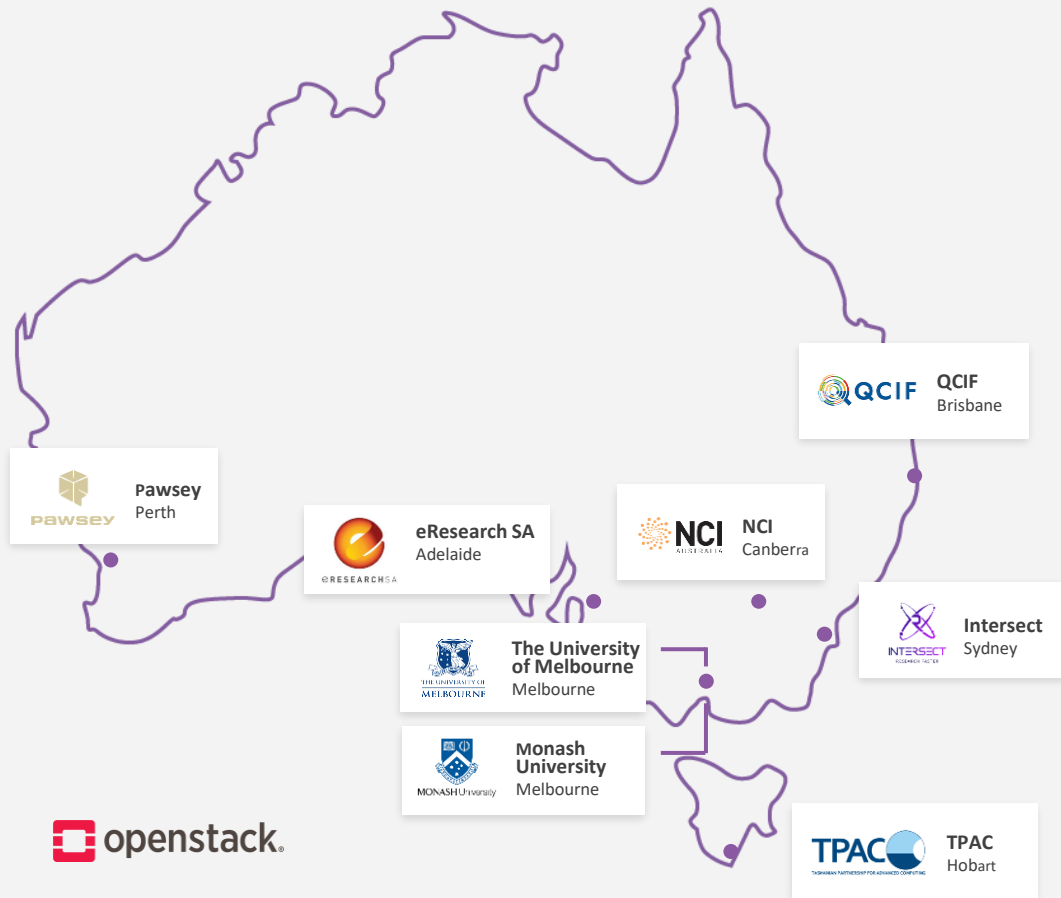
Tech & Ops Workshops, 2020 and 2022

Nectar Federation - an eResearch community



NECTAR CLOUD — THEN

*A national partnership
between 8 institutions
and research
organisations*



NECTAR CLOUD NOW



8,000+

Virtual Machines



34,000

Physical CPU Cores



150,000+

Virtual CPU Cores



320

VGpus (coming soon!)



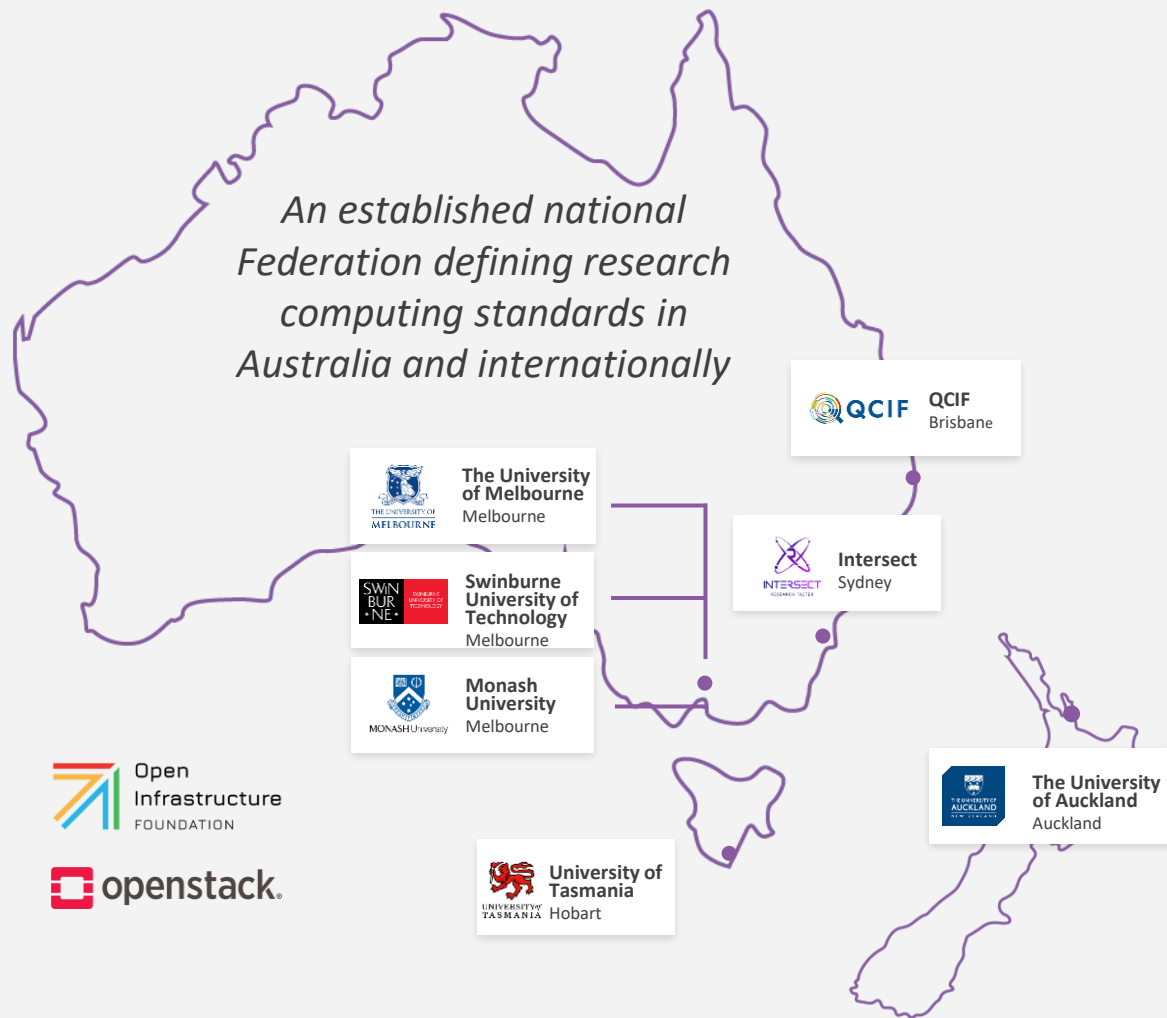
2PB

Object Storage



5PB

File Storage



ENABLING RESEARCH COLLABORATION



LONG-TERM CLOUD COMPUTE SUPPORT FOR RESEARCHERS

Access to local & national datasets



Cost effective for researchers and their institutions



Local and national research cloud compute expertise and support



National research cloud computing standards



Easy to scale resources



Co-Designed

Co-Invested

Co-operated



Easy to build & host platforms due to standard template



Long-term, stable underpinning infrastructure



Multi-compute interoperability



Self service model built for Australian Research



Easier to collaborate across institutions

**SHARED GOVERNANCE
CENTRALISED SKILLS
PROVISION AT SCALE**

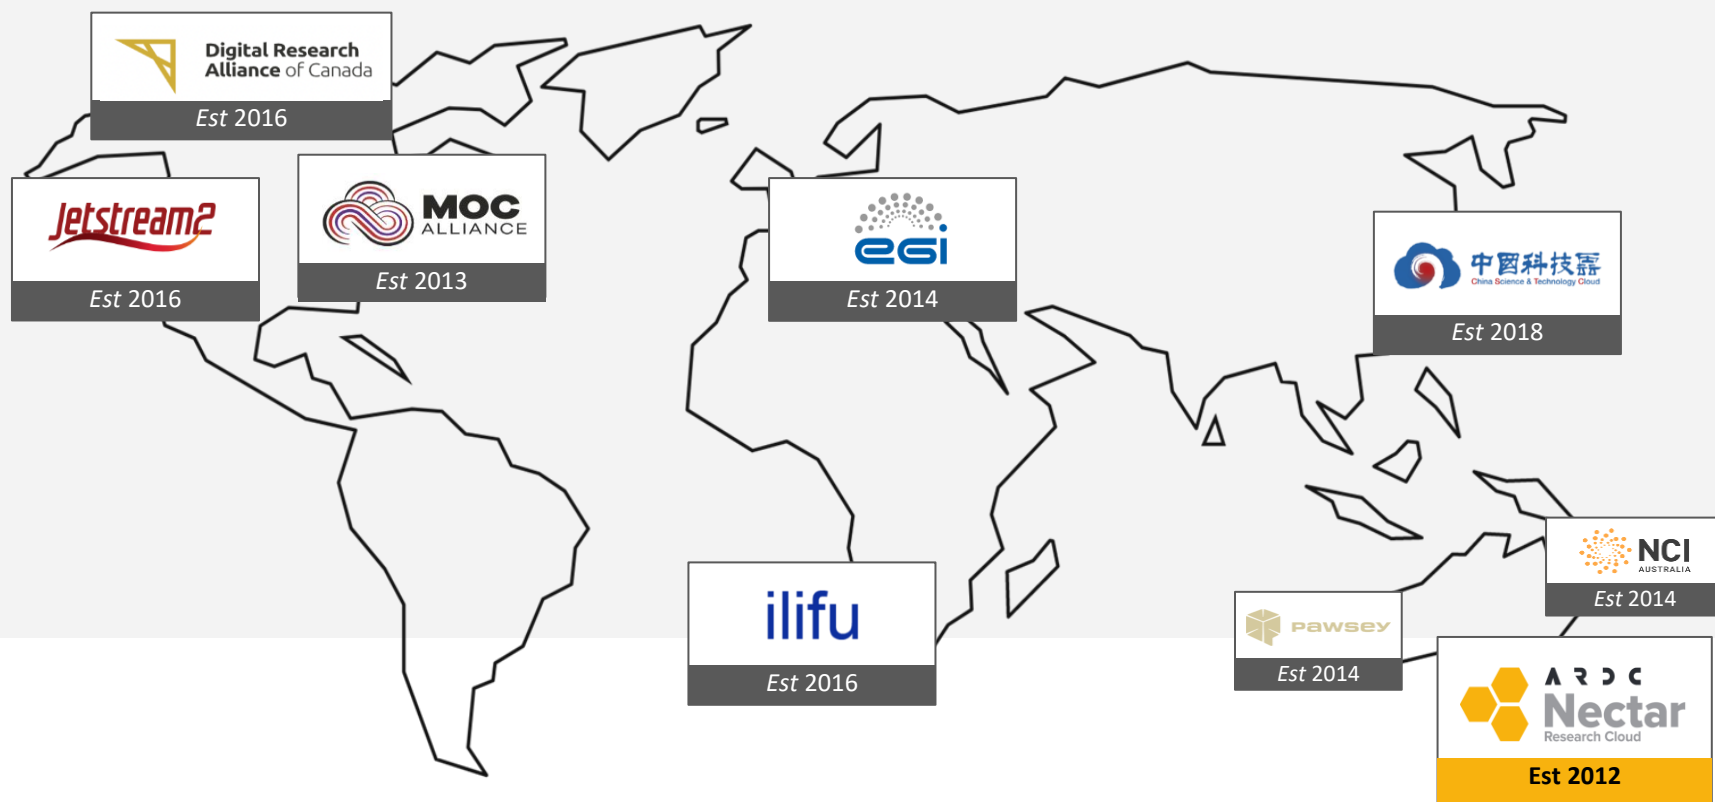
Impact of Nectar



Nectar has been essential to support and grow the computational capability of my research team, which has continued to grow since 2012

Professor James McCaw, *Mathematical Biology Mathematics and Statistics, The University of Melbourne*

FIRST NECTAR THEN OTHER RESEARCH CLOUDS...



SUPPORTING NATIONAL RESEARCH SINCE 2012

In 2021/2022 ARDC Nectar Research Cloud supported:



414
ARC Grants



21
Cooperative Research
Centres and Projects



22,084
Users



3820
Active users



16
NCRIS
Capabilities



27
ARC Centres of Excellence &
Industrial transformation
hubs/training centres



35
Australian
Universities



1,845
Projects



30
Research Platforms /
Virtual Laboratories



5,130
Research
Projects



50,000+
Researchers
Supported annually



251
ARC, NHMRC & MRFF
Research Grants



All
Research Fields



205
NHMRC & MRFF
Grants



2,875
Research Papers



2.3 billion
Virtual-Processor Hours



358
Collaborative Multi-
Institutional Projects

CASE STUDY

Global Research Infrastructure is Key for Fighting the Pandemic



The enhanced Galaxy Australia platform will position Australia at the forefront of bioinformatics infrastructure and substantially improve Australian researchers' access to bioinformatics.

Professor Andrew Lonie,
Director, Australian BioCommons



IMAGE — MOOFUSHI - 390832387 / ADOBE

Biodiversity and Climate Change VL

Virtual Laboratory Case Studies



..decreases the time to complete biodiversity analysis from 2 months to 5 minutes, supporting new applications in research, government and industry

Professor Brendan Mackay,
Director, Griffith Climate Change Response Program

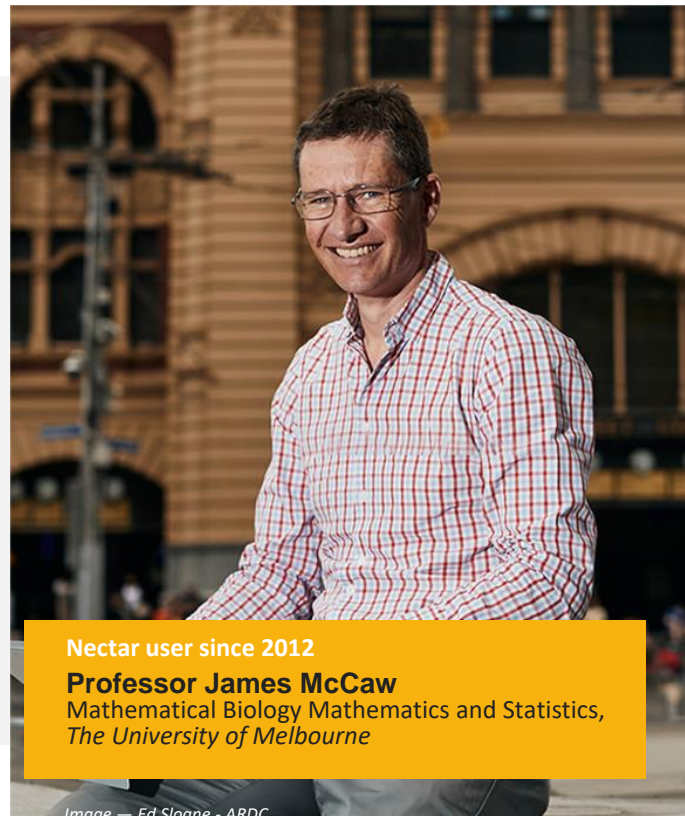


IMAGE — MOOFUSHI - 390832387 / ADOBE STOCK

LONG-TERM CLOUD COMPUTE SUPPORT FOR RESEARCHERS: NATIONAL COVID-19 MODELLING

“

It completely changed my ability to develop my career because the system was merit not dollar-based.



Nectar user since 2012

Professor James McCaw
Mathematical Biology Mathematics and Statistics,
The University of Melbourne

Image — Ed Sloane - ARDC

Nectar Now:

Evolution & Lessons Learned





IS EVOLVING...



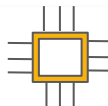
**Jupyter
Notebooks**



**Virtual Desktop
Service**



GPU Service
(248 vGPUs soon!)



**Preemptible
Instances**

Infrastructure as
a Service (IaaS)



**Managed Services
(PaaS)**

Capacity-based
allocation



**Usage-based
allocation (SUs)**

National research
computing standards
(OpenStack)



**Multi-cloud
national standards
(ARCOS)**

Core Cloud
Infrastructure



**High-end
infrastructure**

2022	Restricted Regions Supports Sensitive Data workloads - QCIF pilot & testing
2022	Jupyter Hub National service launched as pilot 18th Oct
2023	Kubernetes Service National managed K8s service developed with use cases from platforms
2023	Machine Learning Service National service

MORE NEW SERVICES!

- Restricted regions QCIF & Monash
- JupyterHub - supported by Core Services
- K8s service- informed by use cases from EcoCommons & ASDC
- Machine Learning will build on ARDC Platform project with UQ, QCIF & Monash initially

Lessons Learned

Building a Community

- New technology and services need a lot of community building, user support, training, promotion, communications, good UI/UX
- Importance of a community of practice (national and international) to share ideas and expertise

Importance of National Investment

- Pioneered virtual labs and Platforms
- Importance of a stable, long-term service
- Importance of merit over dollar-based provision

Managing a Federation

- Has to be a partnership
- Co-design of National standards, policies, processes, governance for interoperability across institutions
- Difficult to attract and retain technical staff
- Trigger points to scale a service nationally

ACKNOWLEDGEMENTS

- Glenn Moloney (Nectar Director)
- Steve Manos, Nigel Ward, Lyle Winton, Paul Coddington (Nectar Cloud)
- Richard Northam, Ian Duncan, Carmel Walsh, Rosie Hicks (ARDC management)
- Hamish Hollewa, Nick Golovochenko, Jo Morris (User Support)
- Distributed helpdesk and User Support staff
- Greg Sauter, Wilfred Brimblecombe, Sengor Kusturica (Core Services Manager)
- Sam Morrison, Andy Botting, Jake Yip (Tech Leads) and the Core Services staff
- Node Directors, Operations Managers, Operations staff
- Nectar, ARDC and Node admin and comms and engagements staff
- Virtual Labs, Platforms, NCRIS facilities, researchers



HOW CAN THE ARDC ACCELERATE YOUR RESEARCH?

Visit us at
eResearch Australasia - Stand 14

2022 Data Driven Research Impact

Download at
ARDC.edu.au

