



Institutional Research Computing Capabilities in Australia: 2024

Meta to Macro

Luc Betbeder-Matibet, UNSW Sydney (luc@unsw.edu.au)

Executive Director Research Technology Services UNSW
Adjunct Associate Professor, Computer Science Engineering, UNSW
President, Australasian eResearch Organisations (AeRO)
Co-chair, eResearch Australasia Conference 2025

Presented at
21st IEEE International Conference on eScience September 18, 2025 Chicago, USA
eResearch Australasia 2025. October 23.

Created by our BOF/Workshop participants @ eRA (2021-24)

Institutional Research Computing Capabilities in Australia: 2024

1st Slava Kitaeff

*The University of New South Wales
Monash University
Melbourne, Australia*

s.kitaeff@unsw.edu.au, slava.kitaeff@monash.edu

2nd Luc Betbeder-Matibet

*The University of New South Wales
Sydney, Australia*

luc@unsw.edu.au

3rd Jake Carroll

*The University of Queensland
Brisbane, Australia*

jake.carroll@uq.edu.au

4th Stephen Giugni

*The University of Melbourne
Melbourne, Australia*

stephen.giugni@unimelb.edu.au

5th David Abramson

*The University of Queensland
Monash University
Brisbane, Australia*

david.abramson@uq.edu.au

6th John Zaitseff

*The University of New South Wales
Sydney, Australia*

j.zaitseff@unsw.edu.au

7th Sarah Walters

*The University of Queensland
Brisbane, Australia*

sarah.walters@uq.edu.au

7th David Powell

*Monash University
Melbourne, Australia*

david.powell@monash.edu

8th Chris Bording

*The University of Western Australia
Perth, Australia*

chris.bording@uwa.edu.au

9th Trung Nguyen

*The Commonwealth Scientific
& Industrial Research Organisation
Melbourne, Australia*

trung.nguyen@csiro.au

10th Angus Macoustra

*The Commonwealth Scientific
& Industrial Research Organisation
Melbourne, Australia*

Angus.Macoustra@csiro.au

11th Fabien Voisin

*The University of Adelaide
Adelaide, Australia*

fabien.voisin@adelaide.edu.au

12th Bowen Chen

*The University of Adelaide
Adelaide, Australia*

bowen.chen@adelaide.edu.au

13th Jarrod Hurley

*Swinburne University
Melbourne, Australia*

jhurley@swin.edu.au

Thanks to the co-authors and especially Slava:

CSIRO – Trung Nguyen, Angus Macoustra

Monash University – David Powell, David Abramson, Slava Kitaeff

Swinburne University – Jarrod Hurley

University of Adelaide – Fabien Voisin, Bowen Chen

University of Melbourne – Stephen Giugni

University of Queensland – Jake Carroll, David Abramson, Sarah Walters

University of Western Australia – Chris Bording

UNSW Sydney – Slava Kitaeff, John Zaitseff



Scope and Reason

- **Our main Academic and National Science Institutions... (those with significant clusters/cores):**
UNSW, Monash, UQ, UMelb, CSIRO, Adelaide, UWA, Sydney, Swinburne
- **A Snapshot:** 2024 (October) leveraging 22-23 BOF work with an improving series planned.
- **Across immediately obvious and simple categories:**
Compute, Storage, Usage, Operations & Strategy
- **In comparison with:**
Each other and with Tier-1 national HPC
(National Computing Infrastructure & Pawsey Supercomputing Centre)
- **As part of an evolving National funding landscape**
and a currently ongoing funding of HPC/Clusters within our institutions.
(We wanted a corpus of data to support our decision in the sector)

Definitions and contexts

Institutional Research Computing in Australia is seen to provide...

- **a capability to support a wide range of research disciplines** (legitimate research instrument for mature and emerging scientific workloads) that is distinct from Tier1/National.
- **main characteristics:**
 - Training & skill development for wide range of users
 - Access for prototyping
 - Data sovereignty compliance

Context: Our National Ecosystem is very small by USA/Europe benchmarks

Australia hosts only 2 national Tier-1 HPCD facilities accessible to researchers :

- NCI Australia and
- Pawsey Supercomputing Centre

With Additional cores in Nectar Australia's Research Cloud

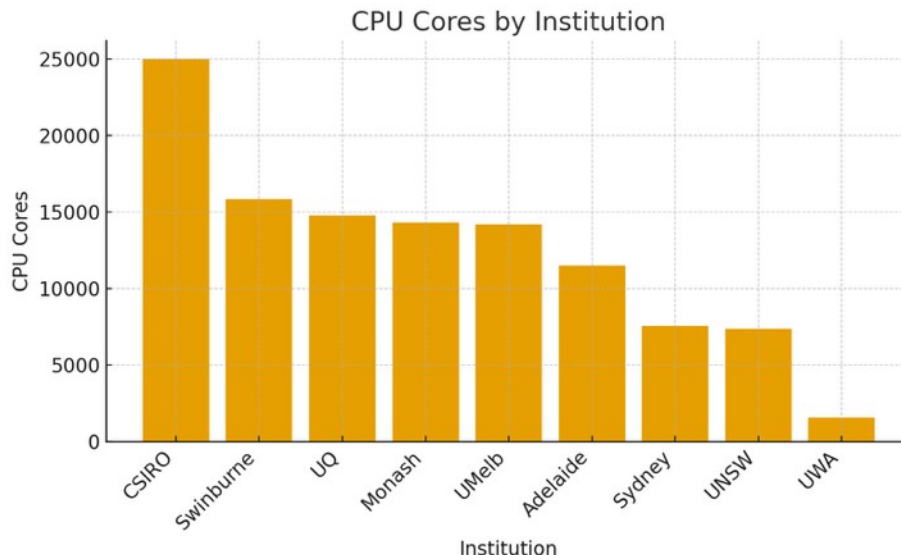
These are primarily funded by our National Collaborative Research Infrastructure Strategy (NCRIS)



High level findings- 2024 data

CPU

- **Institutional CPU core count = ~112K Total**
- **Mixed landscape:** medium-scale systems (with broad or specialised workloads)
- (AMD EPYC adoption alongside Intel)



2024 Australia ~570K CPU cores

Institutional CPU Cores

CSIRO (~25K),
UQ Bunya (~15K),
UWA Kaya (~1.6K)

National CPU Cores (10x to 200x)

NCI (~250K)
Pawsey (~205K)

~21% of national CPU capacity

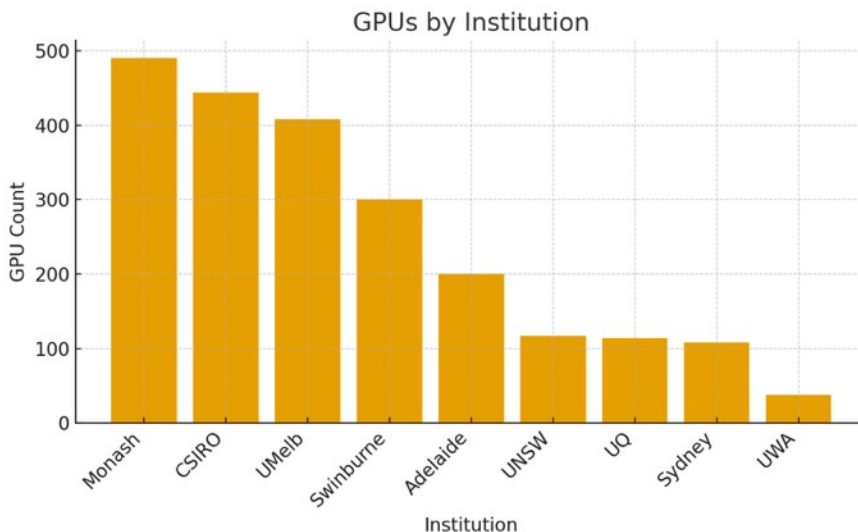
Frontera ... ~450K CPU Cores

Frontier... ~615K CPU Cores

High level findings- 2024 data

GPU

- **Institutional GPU count = 2.2K** (growing fast, 2025 expected double)
- **Currently concentrated:** Monash, CSIRO, UMelb, Swinburne (74%)
- **Generational diversity:** V100, A100, H100, AMD MI210/MI300X
- AI/ML dominates +(chemistry, MD, astronomy workloads)



2024

Institutional GPUs

CSIRO (Virga 444)
Umelb (Spartan 408)
UQ (Bunya 114)

National GPUs (0.5 to 2x)

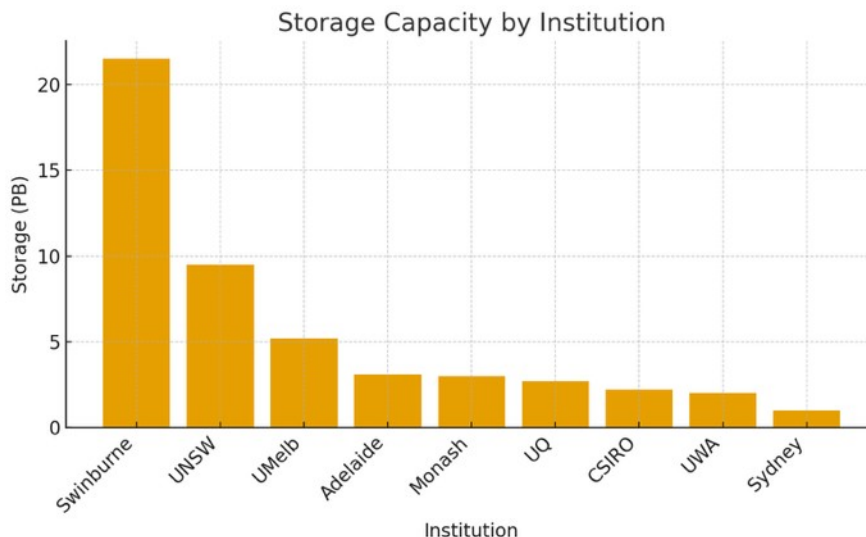
NCI (~200)
Pawsey (~800)

~43% of national CPU capacity,

High level findings- 2024 data

Storage

- **Institutional sector total = 14.8 PB+ scratch** (but many are ~50PB total)
- Mostly <5PB with High utilization (>80%)
- Lustre (CSIRO, Monash, USyd, Swinburne) vs GPFS (UMelb, UoA, UWA)
- Flash tiers emerging (Monash, UQ, Adelaide)
- Significant storage-to-core ratio variation: UNSW 1.29 TB/core vs CSIRO 0.09 TB/core



2024

Institutional Storage (PB)

Swinburne (1.5 Scratch / 20 User)

UQ (2.1 Scratch)

UNSW (10 Scratch/User)

National Storage (2x to 20x)

NCI (~20 Scratch)

Pawsey (~14 Scratch)

High level findings- 2024 data

Alignment with National Facilities

- Institutional (tier 2) computing capabilities have (broadly been) considered by our managers and users to be complimentary with National (tier 1).
 - Development pathway between desktops and national HPC
 - Specialised workflows (sovereignty, interactive, niche needs)
 - Training and community capability building
- Examples:
 - In WA: UWA + CSIRO directly integrate with Pawsey partner shares
 - East Coast: UNSW (\$2m+ pa), Monash, UMelb invest in NCI in addition
- Data storage not integrated but Growing use of Data transfer nodes / Globus (Good quality network)

High level findings- 2024 data

Research Communities & Applications

- **6,000+ active researchers**. Significant national contribution.
- Scale: UWA ~150 → UQ ~1,350
- **Disciplinary diversity:**
 - UNSW: Science (66%) *our Engineering users are on NCI*
 - Monash: Genomics & Medicine
 - UQ: AI/ML & Hypersonics
 - CSIRO: CFD & Climate
 - Swinburne: Astronomy (70%)
- Tools: Python, R, MATLAB, CryoSPARC, Gaussian, GROMACS, Ansys

High level findings- 2024 data

Operational Models & Funding

- **Staffing:** 3–10 FTEs per site
- **Roles:** sysadmins, RSEs, data specialists, community support
- **Funding:**
 - Some central (UQ, Adelaide),
 - Some mixed - Local/central/national (UNSW, UMelb, UWA),
 - Some with large % national / partnerships (Monash, Swinburne)
- **Data centres:** trend to commercial facilities; (on-premise persists at UNSW, UMelb, Adelaide, Swinburne). GPU impact starting to hit.

High level findings- 2024 data

Observations: **Strategic Value of Institutional Infrastructure**

- **GPU ramp up. Speed and cost. Ability to respond.**
 - ~20% of national CPU capacity but ~43% of GPU capacity
- **Good value per user**
 - estimate.. ~\$144M AUD (\$100M USD) Total replacement cost for 6000 Users.

- **Complementary (or is it masking issues in national infrastructure).**
 - Accessible, responsive computing
 - Supports training, methodology development, specialised workflows
 - Builds local expertise and research capability
- **But not strategically aligned.**
 - Acknowledged strategic gap in national planning,
 - Refreshes and planning are organic/local/ ad-hoc

High level findings- 2024 data

Observations: Shared Challenges & Future Directions

- **GPU demand growth > \$ sustainability > placement challenges**
 - Step-change vs sustained demand?
 - Power & cooling constraints > placement and partnering challenges
- **Storage pressures (>80% utilisation)**
 - Step-change vs sustained
 - Classification and sovereignty > certifications
- **Staffing retention and hire risks**
 - Specialists > AI / GPU boom retention and cost
- **Equity of access to resources. New workload requests for GPU**
 - GPU concentration / allocation approaches
 - Lack of national AI systems.
 - Pressure on GPUs in current clusters from new disciplines.

High level findings- 2024 data

Observations: Challenges & Future Directions

- **Future directions (BOF work)**
 - Standardised metrics & reporting
 - Shared training & documentation

- **Opportunities (partner work)**
 - Federated identity management? National program
 - Data alignment / certification for sensitive workloads
 - Strategic funding and differentiation planning
 - Data Centers ... placement efficiency?

- **Risks**
 - Dependency on national funding and programs (timing and scale).

High level findings- 2024 data

Summary:

- **Institutional Computing is an acknowledged core component of our national HPCD research infrastructure mix.**
 - 20% of CPU and 43% of GPU capacity
 - Supports 6,000+ researchers across disciplines for "only" \$150m
- **We are at an Interesting stage in our national and institutional planning**
 - Sharing our numbers has helped our local business cases and co-investments nationally.

Takeaway:

Institutional computing bridges everyday research and national-scale discovery

Sharing + Publishing system and service data for sector benchmarking is helping us.





Questions

luc@unsw.edu.au