



National Computational Infrastructure

Australian National University

SCA25 APRP, SKA and Genomics

Andrew Howard – Associate Director (Cloud Services & Innovation)

March 2025

We acknowledge and celebrate the Traditional Owners on whose lands we meet, and pay our respect to the Elders past, present and future.



Artist: Lynnice Letty Church **Tribes:** Ngunnawal, Wiradjuri & Kamilaroi (ACT and NSW)
Gadi - "to search for" in Ngunnawal language - January 2020 for NCI Gadi Supercomputer





Intro

Agenda

Agenda

- Introduction
- What is NCI
- NCI Updates
 - New Director
 - Nirin Cloud hardware refresh
 - Australian Copernicus Hub upgrade
 - Australian Intelligence
 - APRP
 - Foundational National Genomics Services







National High Performance Compute + Data

New NCI Director

National Computational Infrastructure

Staff Update

- Former director Dr Sean Smith left ANU and NCI December 2024.
- Prof Andrew Rohl will commence as NCI Director 24-Mar-2025.



National Computational Infrastructure

Staff Updates

- NCI welcomes our new interim director Prof Andrew Rohl and Senior Operations Manager (interim) Dr Hanah Rodriguez Colombo
 - Andrew had had a long association with HPC in Australia and established the Pawsey Supercomputing Centre in 2013. He has extensive leadership experience including setting up a university-wide data science centre at Curtin as well as leading the largest school in Science and Engineering, and will commence at NCI on the 24th of March 2025.



National Computational Infrastructure

Staff Updates

- We also welcome our new interim Senior Operations Manager Dr Hanah Rodriguez Colombo
 - Hanah is currently the EO to the Chief, Research Information Services at ANU and brings a wealth of experience in research management and operational management of large complex organisations and will commence on the 19th of March.



National Computational Infrastructure

Staff Updates

- I would like to thank Prof Ute Ute Roessner Pro-Vice Chancellor, Research Initiatives and Infrastructure, Research and Innovation Portfolio for her role as Acting Director of NCI over the past year.
- Prof Roessner has been appointed as the CEO of the Australian Research Council and will commence in April.





National High Performance Compute + Data

What is NCI?

National Computational Infrastructure

Overview

- NCI is Australia's National High-Performance Computing service
 - comprehensive, vertically-integrated research service, providing national access on priority and merit
 - driven by research objectives
- Operates as a formal collaboration of ANU, CSIRO, the Australian Bureau of Meteorology and Geoscience Australia (foundation collaborators).
- As a partnership with research-intensive Universities, supported by the Australian Research Council
- NCI is funded directly by the Australian Government, through the Department of Education, Skills and Employment's National Collaborative Research Infrastructure Strategy (NCRIS).




Where are we located?



Spectrum of Science

FUNDAMENTAL



- Physics
- Chemistry
- Mathematics
- Astronomy

STRATEGIC




- Environment
- Medical
- Geoscience
- Agriculture
- Materials

APPLIED



- Weather forecasting
- Extreme weather
- Disaster management and mitigation

INDUSTRY



- Hydrological modelling
- Medical research institutes



HPC & Storage

Our Systems

Computational Systems - Gadi

- Local Ngunnawal indigenous word 'Search for'
- Commissioned Nov 2019 / Jan 2020
- HDR InfiniBand
- 3024 Compute (CLX) Nodes
 - 145,152 Cascade Lake Xeon Cores (Platinum 8274, 2.9-4.0 Ghz/24c)
- 720 Compute (SPR) Nodes
 - 74,880 Sapphire Rapids Xeon Cores (Platinum 8470Q, 2.10-3.80 Ghz/52c)
- 54 Large Memory Compute (CLX) Nodes
 - 2,592 Cascade Lake Xeon Cores
- 804 Compute (BDW) Nodes
 - 22,512 Broadwell Xeon Cores
- 160 GPU (Volta) Nodes
 - 7,680 Cores, 640 V100 GPUs
- + other large mem, DGXs, special purpose nodes

- 4968 nodes, 259,560 cores total



'Computation is only there to extract the value of the data'

- NCI data intensive user




Nirin

Cloud Systems

Nirin Cloud

- NIRIN is the term for “Edge” in the First Nations Wiradjuri language.
- NCI’s NIRIN Cloud is an Infrastructure as a Service (IaaS) capability based on OpenStack which allows us to deliver the “edge case” services which fit alongside and support our traditional HPC workloads.
- NIRIN is used by both NCI and our stakeholders to deliver long running and web facing services.



NCI Cloud

World-class, Dynamic, Cloud Capabilities

- Repurposed RAIJIN (Australia's formerly fastest research Supercomputer) into NCI Cloud
 - OpenStack based
 - 22 racks
 - Distributed across two data centres
 - ~1,440 Intel Xeon Broadwell cores
 - ~42,000 Intel Xeon Sandy Bridge cores
 - 200 Tb of memory
 - 4 * Intel Optane large memory nodes 3.6Tb memory
 - Integrating K80 GPUs
 - 6 petabytes of CEPH operational disk storage
 - Volumes
 - CephFS
 - S3
 - 100G & 56G Ethernet interconnect
 - Dispatchable compute
 - (power up and down based on demand)



Nirin – Supporting NCI Service Delivery

- NIRIN supports the delivery of NCI's mission and provides the platform for:
 - Public facing services
 - Data Collections
 - Data analytics and dashboards
 - Data inflow and processing
 - Service portals
 - Virtual Laboratories and Research Environments



Nirin – Supporting NCI Stakeholders

- NIRIN is also the delivery platform for many of NCI's Stakeholders Cloud based applications, services and workflows over the full spectrum of the science disciplines NCI supports.
- Infrastructure as a Service allows our stakeholders to configure the collection of programmable infrastructure Compute, Storage, Networks and Security rules which can be scaled up and down to meet their requirements.



Nirin – Hardware Refresh

- We have extended the life of the old HPC hardware for over five years delivering a significant benefit for Australian Tax payers.
- The hardware is beginning to fail and does not meet our users performance requirements.
- A\$1.2M allocated to replace the majority of the hardware.
- Approach to market completed and vendor selected.
- Delivery date TBC



National Computational Infrastructure

NDRI Australian Government Funding round

- National Digital Research Infrastructure (NDRI) refers to the digital tools and services that are collectively managed and operated for researchers to access and collaborate across the country.
- The Australian Department of Education has identified A\$400M for support of various National Collaborative Research Infrastructure Strategy (NCRIS) programs.
- NCI and Pawsey have submitted proposals for Infrastructure refresh for both facilities to replace/augment our existing HPC and Storage capabilities.
- Expected announcement before the next Australian Federal Government election (tentatively) expected to be announced in coming weeks.





pawsey



NCI
AUSTRALIA



Australian Research Data Commons



aarnet

Australia's Academic
and Research Network



DIGITAL NEXUS
ALLIANCE

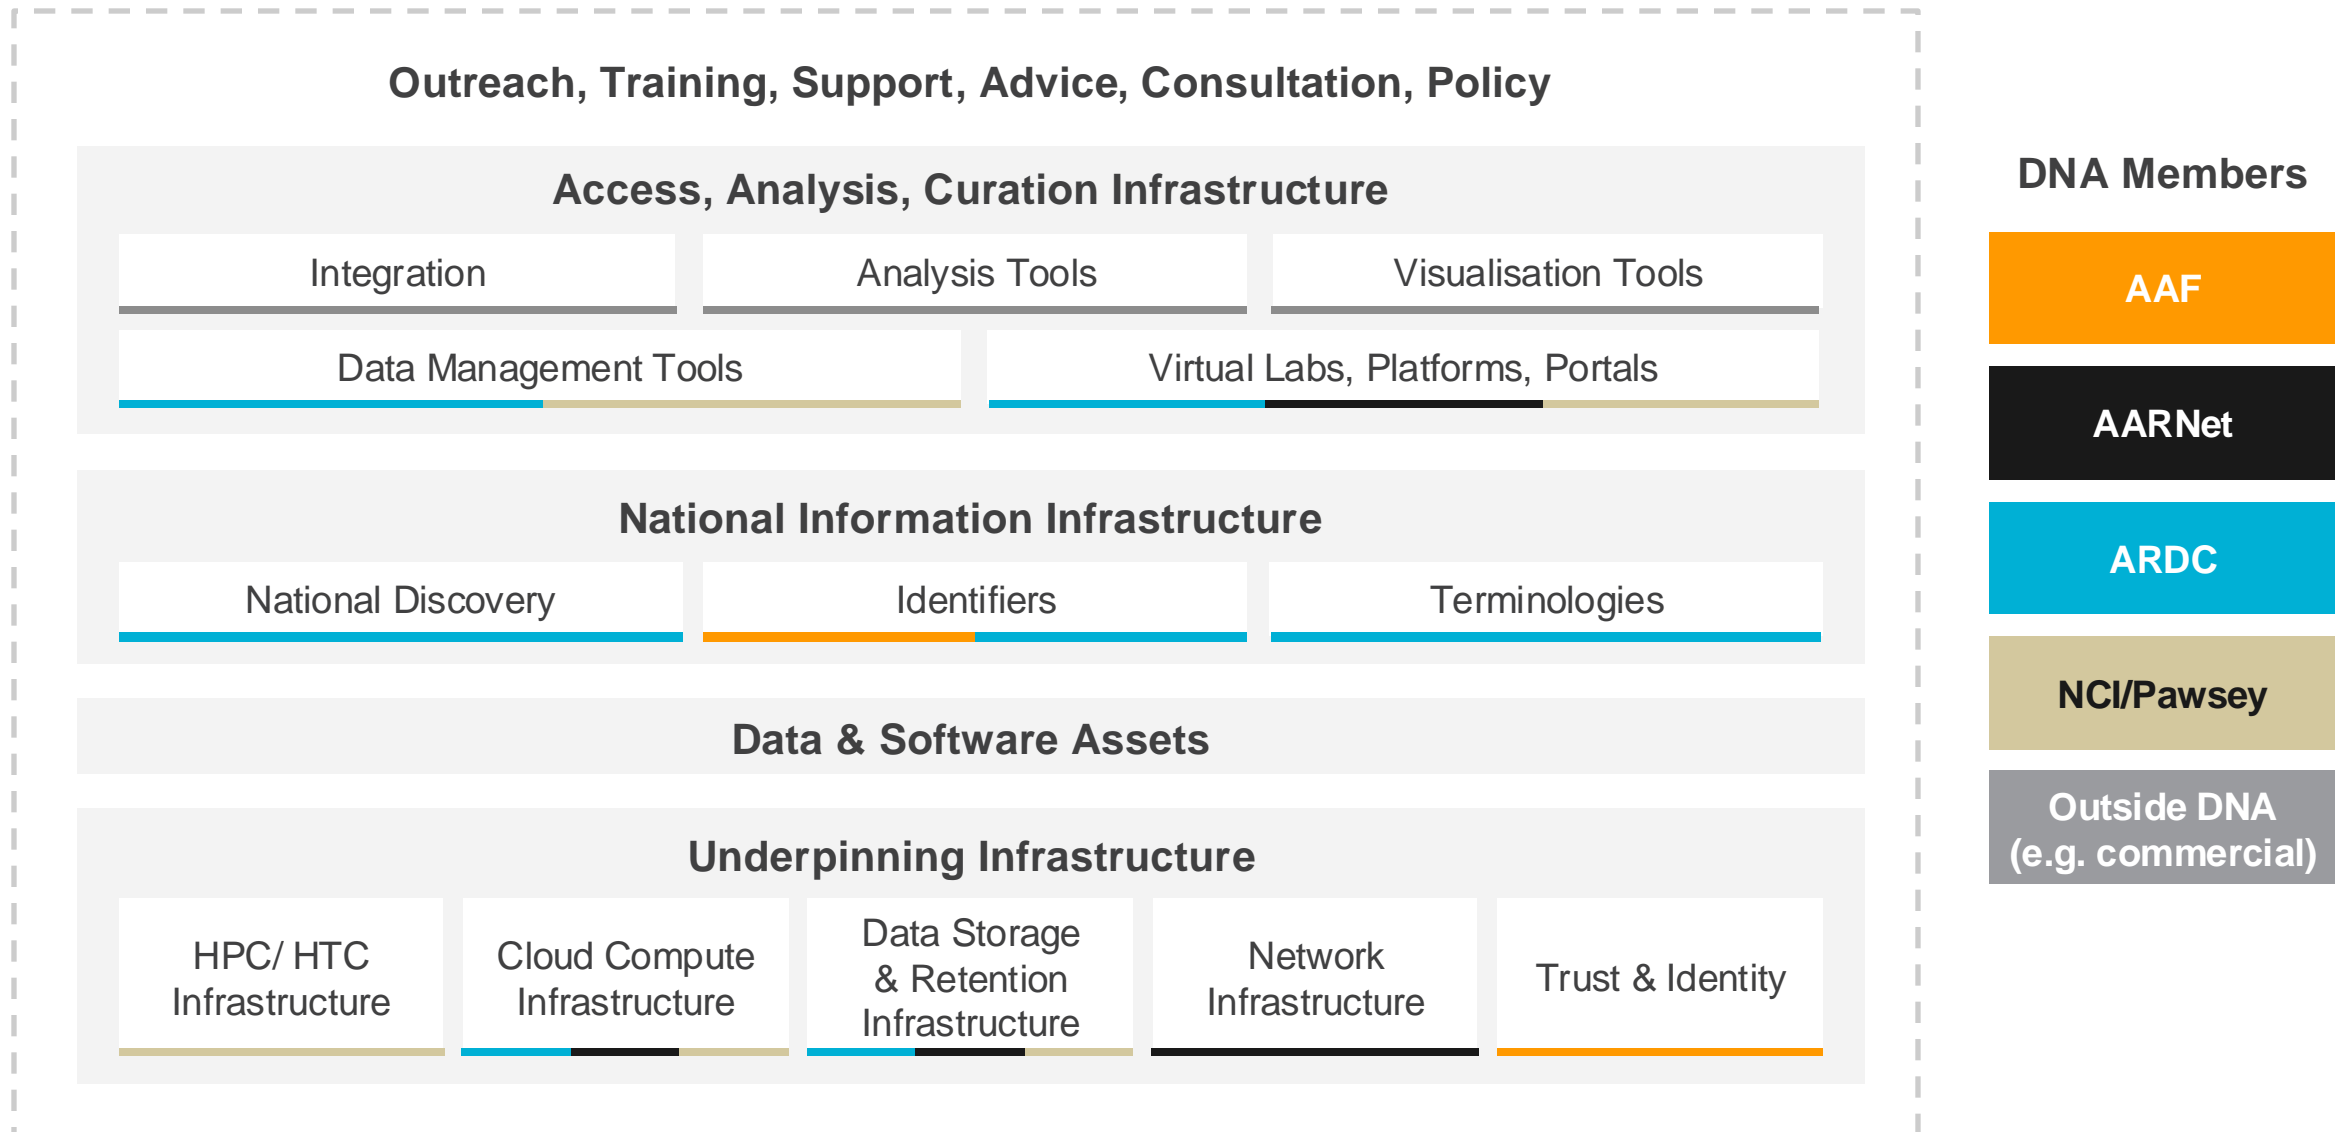


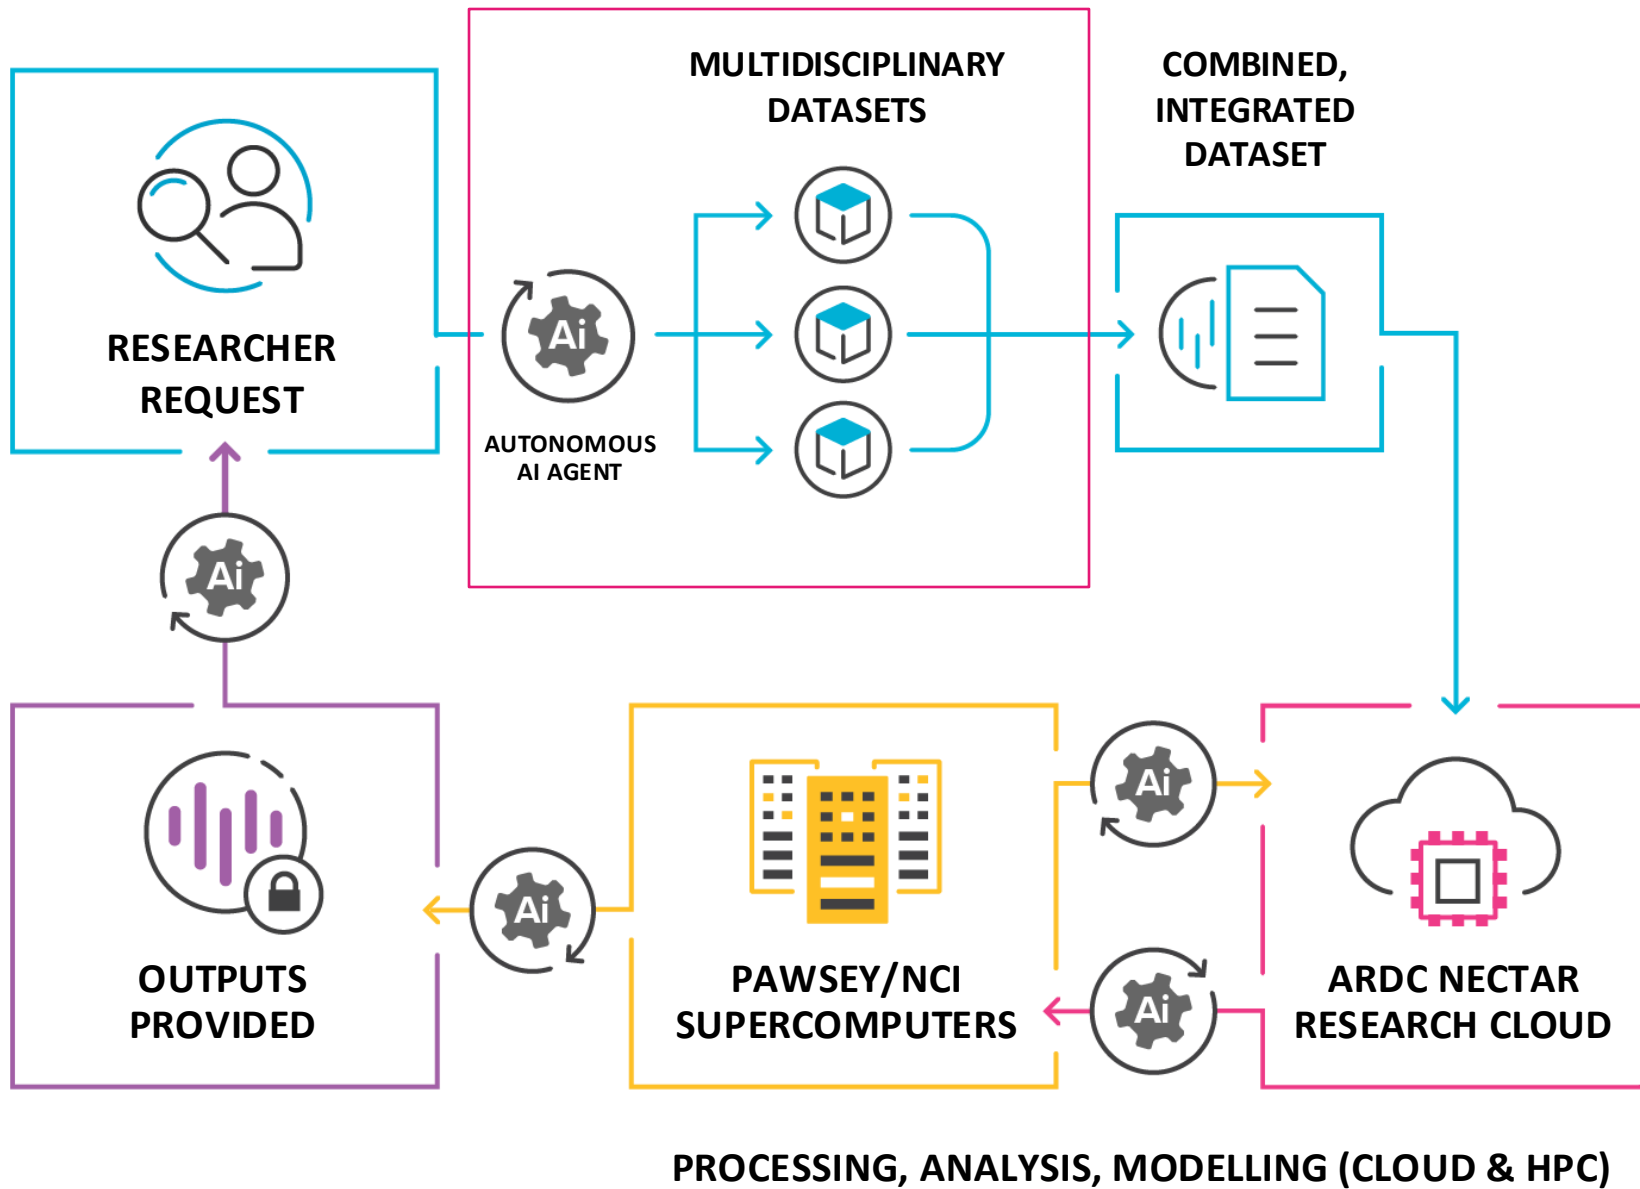
AUSTRALIAN
ACCESS FEDERATION

Digital Nexus Alliance (DNA)

- The Digital Nexus Alliance DNA refers to the collective foundational capabilities for Australian national research enabled by the ongoing collaboration between
 - Australian Research Data Commons (ARDC)
 - Australian Academic and Research Network (AARNet)
 - Australian Access Federation (AAF)
 - National Computational Infrastructure (NCI)
 - The Pawsey Supercomputing Centre

DIGITAL NEXUS ALLIANCE





NDRI OPPORTUNITIES FOR ALL OF US



Complement



Collaborate



Connect



AusCopHub

Australian Copernicus Hub upgrade

Area of interest < 2025



Area of interest > 2025

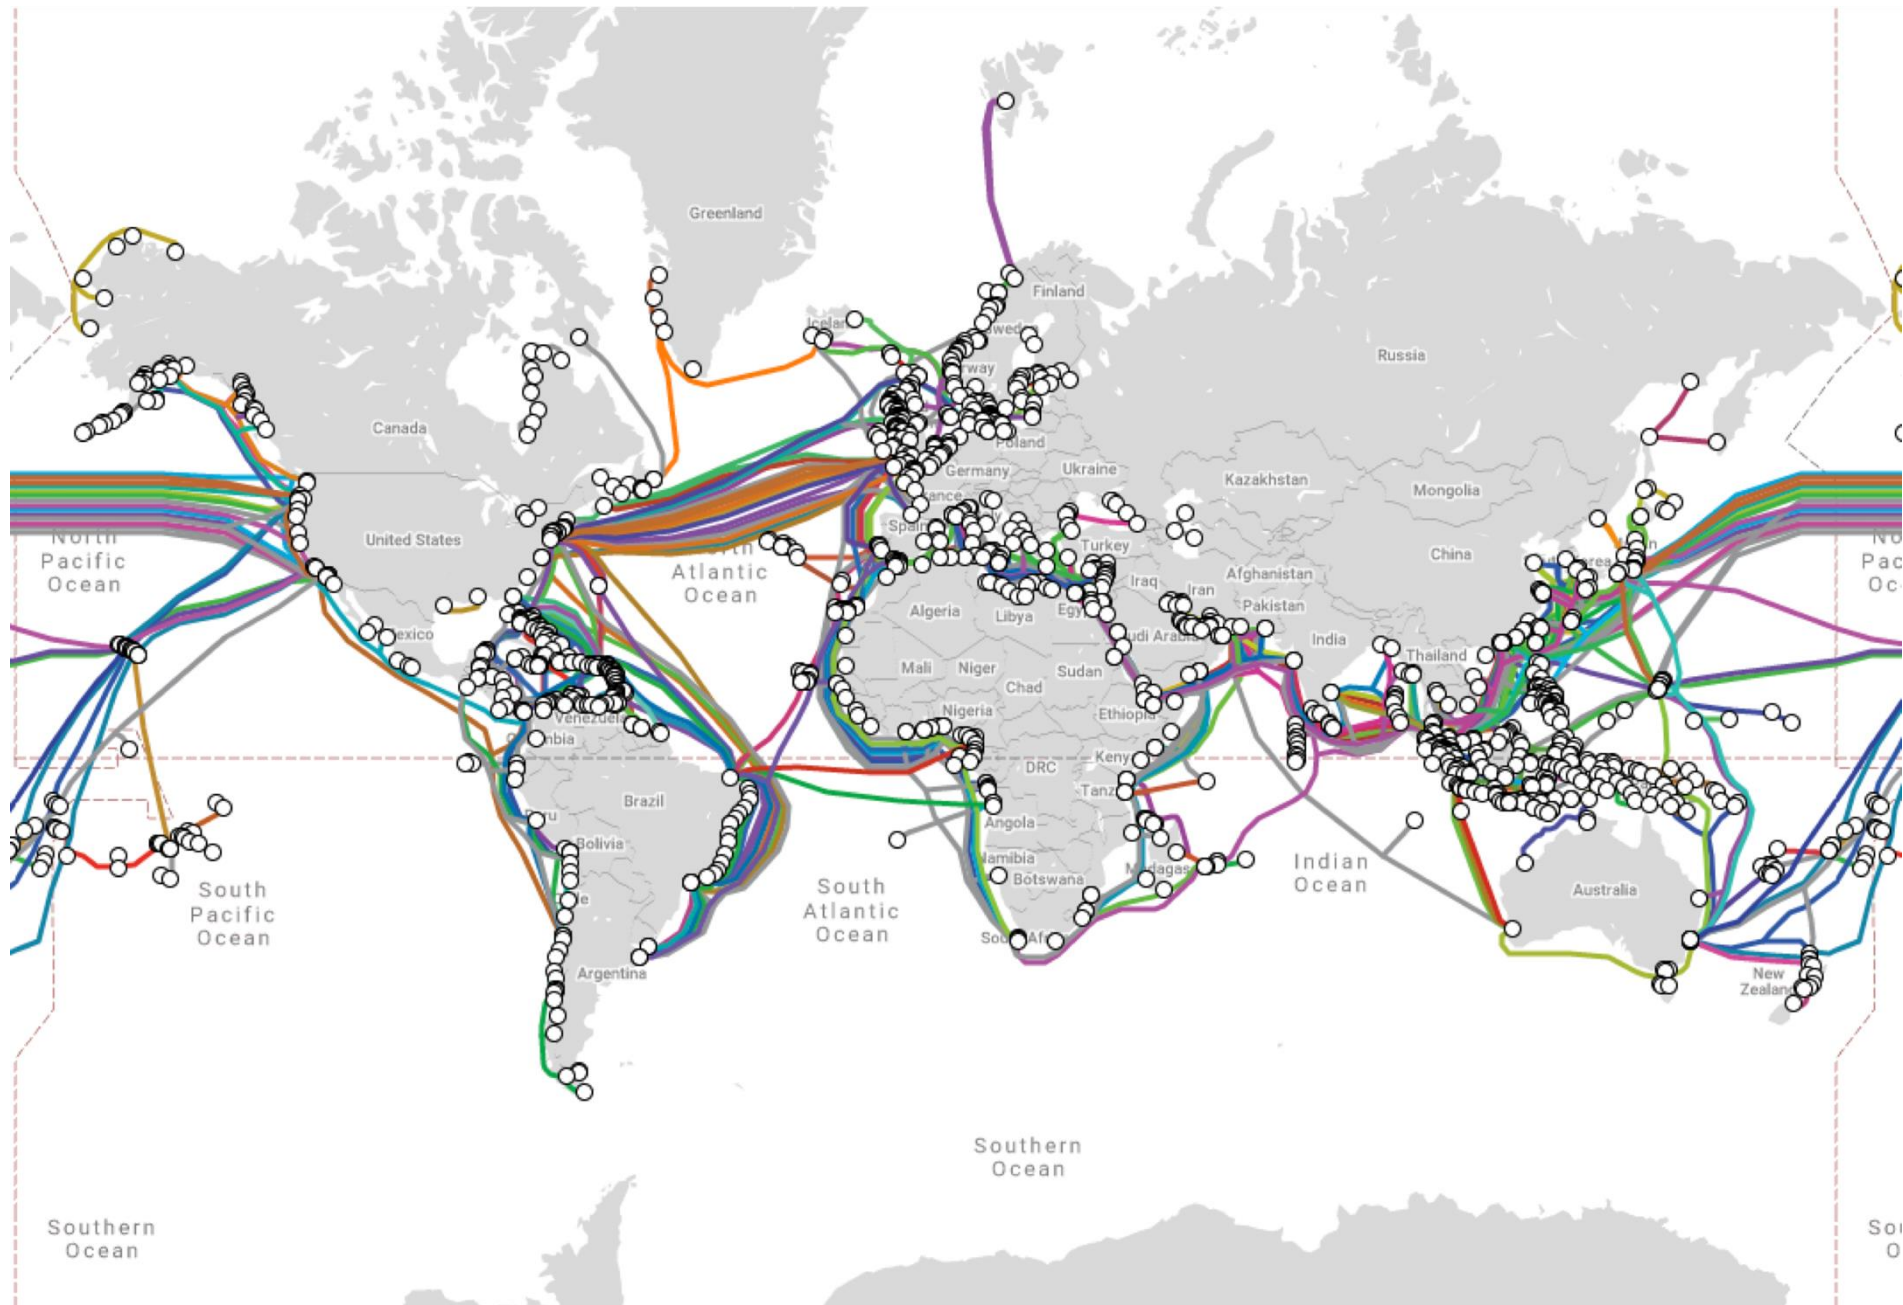


Upgrade for current ESA standards (CDSE)

- Copernicus Data Space Ecosystem (CDSE)
- AusCopHub is operated by NCI on behalf of Australian Government Department: Geoscience Australia.
- AusCopHub has been in operation for around 12 years.
 - Supporting software developed by my team around 12 years ago.
- Over 12 years we have accumulated 5Pb of Sentinel Earth Observation Data.
 - Cost of operation around A\$1M per annum.
 - Delivers A\$5 Billion in GDP contribution with projected increase in GDP contribution in 2030 to A\$30 Billion.
- New ESA standards (CDSE) introduced late 2023 with planned introduction 2024.
- Australia is the first nation to deliver a service based on the new standard.

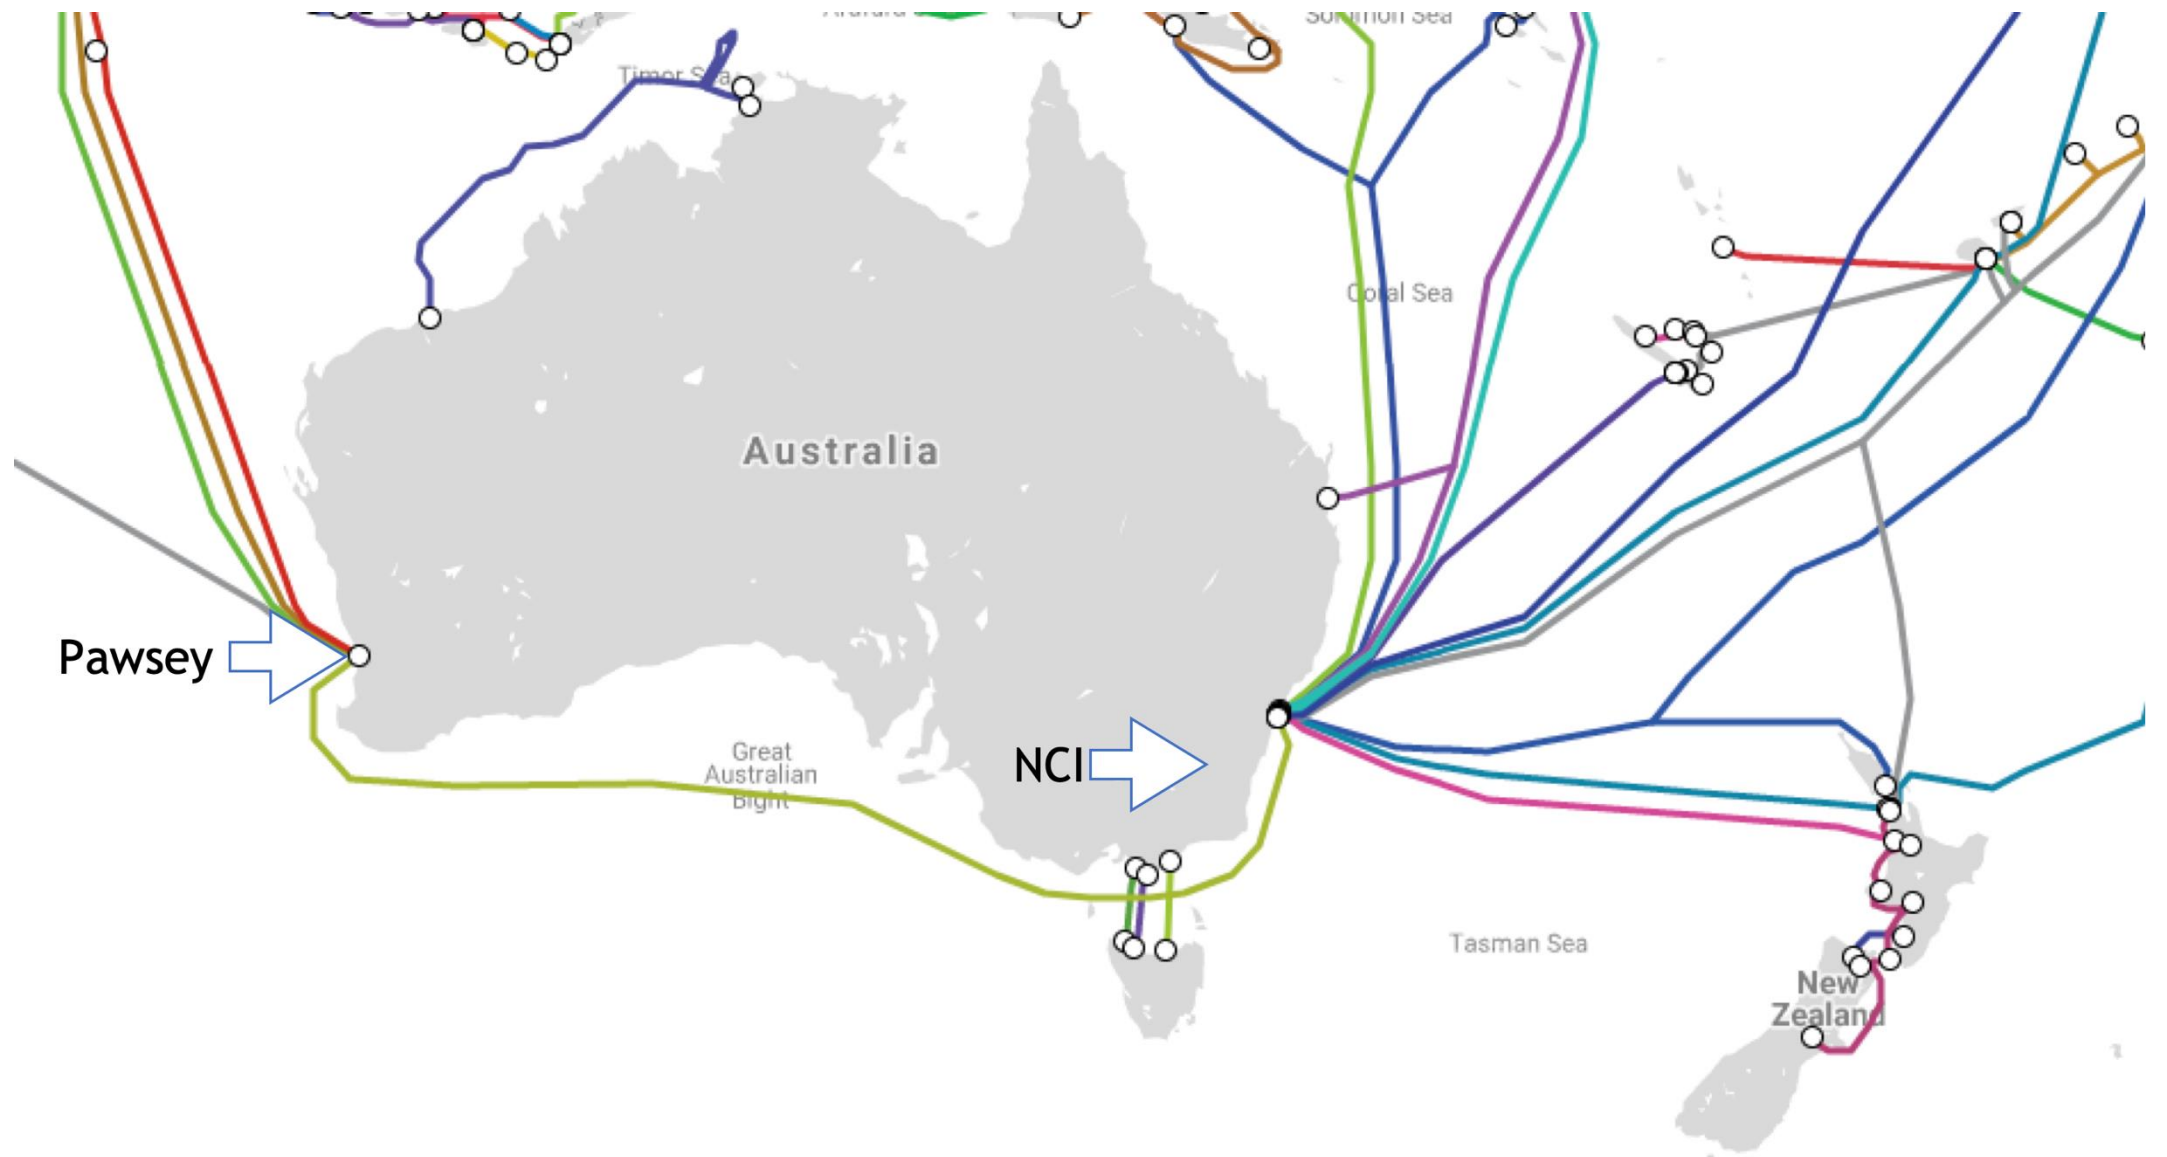
CDSE

- <https://dataspace.copernicus.eu/>




Network tuning

- Based on the groundbreaking work done by our ESNNet colleagues on fasterdata.es.net
- Australia geographically presents the most challenging and extended global R&E data paths.
- NCI worked closely with the CDSE hubs to tune for the fastest possible access over Long Fat Pipes.
- In collaboration with AARNNet
 - Failover testing for alternate paths from Australia to Europe will occur next week.




A wave of clouds Over the Gulf of Carpentaria.

- Morning glory in Queensland (<https://dataspace.copernicus.eu/gallery/2024-10-5-morning-glory-queensland>)



Ball's pyramid in Australia

- Following shadows with Sentinel-2
- <https://dataspace.copernicus.eu/gallery/2024-2-13-balls-pyramid-australia>
- February 29th, a leap day that occurs in the calendar only once every four years, shows how much we humans struggle to grasp the concept of time.



AusCopHub production status

- In recent weeks in collaboration with the French development company Gael Space Systems we have increased download speed by 10x delivering **16Tb** per day.
- Beta release to State Government agencies occurred 1-Mar-2025.
- Full production release 1-Apr-2025.
- Old data will continue to be available in Read only mode with new products from upgraded hub from 1-Feb-2025.
- Old data will be progressively integrated into the new hub architecture over 2025.
- Expand region of interest to include complete Antarctic coverage.

Other Hubs in APAC and Digital Twins

- Phillipines Copernicus Hub installation in progress. See presentation in Disaster Management session on Monday.
- Digital Earth
 - Antarctica
 - Australia
 - Africa
 - Pacific

Australian Intelligence (AI)

- Yesterday at SCA25 we announced **Australian Intelligence (AI)**.
 - <https://www.nci.org.au/news-events/news/nci-unveils-australian-intelligence-ai-a-revolution-data-analytics-and-ai-powered>
- **Australian Intelligence** is designed to empower researchers by providing a seamless integration of advanced analytics and AI-driven solutions. Built on a foundation of robust cybersecurity measures, this service ensures that sensitive data remains protected while still enabling powerful insights and decision-making capabilities.
- **AI** delivers a set of curated AI/ML models integrated into the NCI Australian Research Environments (ARE).... *ARE is based on the OSU (Open on Demand) platform.*

Genomics Workflows and Services

- NCI is delivering the foundation capabilities for Australia's National Genomic services for BioPlatforms Australia,
- Data Portals, Storage and Cohort selection
 - Gen3
 - LocalEGA
- Trusted Research Environments
 - Access to tools to analyse genome data with only the Analysis products exportable
- Tools and utilities
 - Based on EMBL-EBI
 - GA4GH
 - Containerised services (HPC and Cloud compatible)
 - Ready to use genomics research environments

SKA (Square Kilometre Array)



SKA (Square Kilometre Array)

- The SKA project is an international effort to build the world's largest and most capable radio astronomy observatory, designed to enable transformational science that will change our understanding of the Universe.
- Construction is underway of the Australian components of SKA.
- Australia will host the SKA Observatory's low frequency telescope (SKA-Low).
 - SKA-Low will comprise an array of 131,072 'Christmas tree-shaped' antennas, grouped in 512 stations, each with 256 antennas. The antenna stations will span out along three spiral arms, stretching 65 kilometres end to end.
 - SKA-Low is located at CSIRO's Murchison Radio-astronomy Observatory, a legislated radio quiet zone in remote Western Australia, around 800kms north of the state's capital city, Perth, and 315kms north east of regional centre Geraldton.
 - Local indigenous community consulted and closely involved in the location, construction and ongoing maintenance of the facility.

Conclusion

- All Sentinel data is freely available (thank you ESA).
- Australian State Government agencies have switched to new hub, full public release 1 - Mar-2025
- AusCopHub provides faster access for APAC based users.
 - Philippines hub currently being installed and tested
- Packaged Genomics analysis and data management toolkits for use by the Australian Research community.



Questions ?



NCI Contacts



General enquiries: +61 2 6125 9800




Support: help@nci.org.au



Email: andrew.howard@anu.edu.au

Email: daniel.rodwell@anu.edu.au



Address

NCI, ANU Building 143
143 Ward Road
The Australian National University

Canberra ACT 2601