# NCI AUSTRALIA

## PORTFOLIO OF SERVICES

NCI Australia
The Australian National University
143 Ward Road
Acton ACT 2601

T +61 2 6125 9800
E enquiries@nci.org.au
W nci.org.au

@NCINews
National Computational Infrastructure
NCI Australia



More information can be found on NCI's website at nci.org.au.
To read this Portfolio of Services online, scan this code:

Produced by
NCI Australia
The Australian National University
143 Ward Road
Acton ACT 2601

Cover image: Turbulent particle motions from high-resolution simulations of star formation.
Image by James Beattie, The Australian National University

We acknowledge the Traditional Custodians of the ACT, the Ngunnawal people.
We acknowledge and respect their continuing culture and the contribution they make to the life of this city and this region.

# CONTENTS



# CONTACT US

NCI Australia provides access to computing, data and virtual environments through schemes allowing access to researchers from many different institutional and discipline backgrounds.

Access to NCI facilities and services by commerce and industry is welcomed. Access is available on a Fee-for-Service basis. If your organisation is interested in getting access to NCI Australia capabilities, please submit a business enquiry at: enquiries@nci.org.au

Documentation about using and accessing supercomputer, cloud and data resources can be found in our online User Guides, visit nci.org.au.

NCI Australia

# ABOUT **NCI**

NCI is a collaborative hub at the centre of the Australian research ecosystem. We enable transformative science through big data and computing technologies, platforms and expertise. Through Extreme Computing and Data Analytics infrastructure – high-performance computing (HPC) and data (HPD) systems – tightly coupled with our internationally renowned expertise, NCI provides essential services that support the needs of research and industry today and into the future.

NCI supports the work of over 5,500 researchers from 35 universities, 5 national science agencies, a dozen NCRIS science infrastructure providers, 3 medical research institutes, and industry. Scientific research is highly dependent on the fusion of “big compute” and “big data” that NCI provides in areas such as weather and climate science, the earth sciences, earth observation, bioinformatics, and astronomy.

*NCI is a collaborative hub at the centre of the Australian research ecosystem. We enable transformative science through big data and computing technologies, platforms and expertise.*

Image: NCI data centre at The Australian National University.

# WHAT WE DO

**NCI Australia is the country's leading high-performance data, storage and computing organisation, providing expert services to benefit all domains of science, government and industry.**

---

The Gadi supercomputer, operated by NCI, can perform over 15 quadrillion calculations per second.

NCI operates multiple high-performance filesystems, and is home to more than 100 petabytes of critical research data.

The Nirin cloud supports interactive data analysis workflows with dedicated environments optimised for specific research communities.

---

NCI's data services allow users, data portals and external science cloud environments to access, interact with and extract value from our climate, weather, astronomy and genomics data collections.

NCI curates and optimises nationally and internationally significant datasets such as CMIP6 and ERA5, making them suitable for data-intensive science as well as broader access.

NCI's VizLab team generate images, videos, and VR experiences that extend the discovery process for scientists, and effectively communicate high-impact research to scientists and the general public.

---

NCI produces training resources and in-person training courses throughout the year to help develop the skills of the HPC, AI and Data Science communities.

NCI provides User Support to assist users with administrative and technical issues around accessing and using systems and services at NCI.

# HIGH-PERFORMANCE COMPUTING
# **GADI SUPERCOMPUTER**

NCI's Tier-1 supercomputer, Gadi, is capable of over 1.5 billion hours of computing per year.

The name 'Gadi' means 'to search for' in the Ngunnawal language. It perfectly encompasses NCI's mission of scientific research and high-performance computing: to search for knowledge that can make the world a better place through innovative world-class research.

- 15+ petaFLOPS (over 15 quadrillion calculations per second)
- 250,000 CPU cores (Intel Sapphire Rapids, Cascade Lake, Skylake, Broadwell) in 4,800 nodes
- 640 NVIDIA V100 GPUs in 160 nodes, 2 NVIDIA DGX A100 nodes
- Linked with storage systems via 200Gb/s HDR InfiniBand by NVIDIA
- More than 200 scientific software packages


Image: Gadi supercomputer, NCI.

# CLOUD COMPUTING
# NIRIN CLOUD

NCI's newest cloud computing platform is called Nirin, meaning 'edge' in the Australian First Nations Wiradjuri language. The Nirin cloud provides a high-availability and a high-capacity zone, closely integrated with the Gadi supercomputer and NCI's multi-Petabyte national research data collections.

The cloud comprises a mix of Intel Broadwell and Sandy Bridge processors and NVIDIA K80 GPUs. NCI has repurposed hardware from the previous Raijin supercomputer to provide powerful, new capabilities that enable interactive data analysis and data preparation processes.

The Nirin cloud platform is tightly integrated with NCI's supercomputing and high-performance storage infrastructure. Located entirely within the NCI system, this provides privileged, high-speed access to the 100 Petabyte global filesystems on site. It also makes Nirin the ideal home for data-intensive tools and services such as data analysis environments, Virtual Laboratories and internal data publishing processes.

Nirin is divided into two zones:

High-Availability Zone

- 1,856 Intel Xeon Broadwell cores
- 22TB memory

High-Capacity Zone

- 16,640 Intel Xeon Sandy Bridge cores
- 32TB memory
- 40 Nvidia K80 GPUs

Image: Nirin artwork by Anthony Best, Canberra-based Indigenous artist. The artwork shows the lines of communication meeting at gathering places, represented by the circular features. The arch shapes represent scientific communities researching in the cloud, surrounded by data.

# DATA STORAGE
# DISK, TAPE & SSD

NCI operates multiple high-speed, low-latency global parallel filesystems to support the most data-intensive research workloads.

Image: NCI filesystems.

We run multiple high-performance filesystems, linking high-performance computing (HPC) with high-performance data (HPD) via multiple 200-gigabit network links. NCI's filesystems, catering to the needs of our research community, enable the next generation of computational tasks, including high-throughput computing.

NCI's filesystems contain around 100 Petabytes of data storage capacity, providing space for research data to be stored in five separate global Lustre filesystems, reaching a total aggregate IO performance of 500 GB/second.

NCI also runs a special IO-Intensive Platform, a dedicated filesystem using 576 4-Terabyte NVMe drives for cumulative performance around 960 Gigabytes per second.

NCI also stores 50 Petabytes of archival project data in state of the art magnetic tape libraries. In total, NCI runs more than 15,000 hard drives from vendors including NetApp, DDN and HPE.

Image: NCI tape library.

# DATA SERVICES & **COLLECTION MANAGEMENT**

NCI increases the value and reach of research-ready datasets by providing high quality and high performance data-intensive services. As a trusted national data repository, NCI optimises some of the largest and most significant datasets requiring our high-performance capabilities. These cover climate and weather collections through CMIP6 and ERA 5, earth-observation through Landsat and Copernicus, and many more.

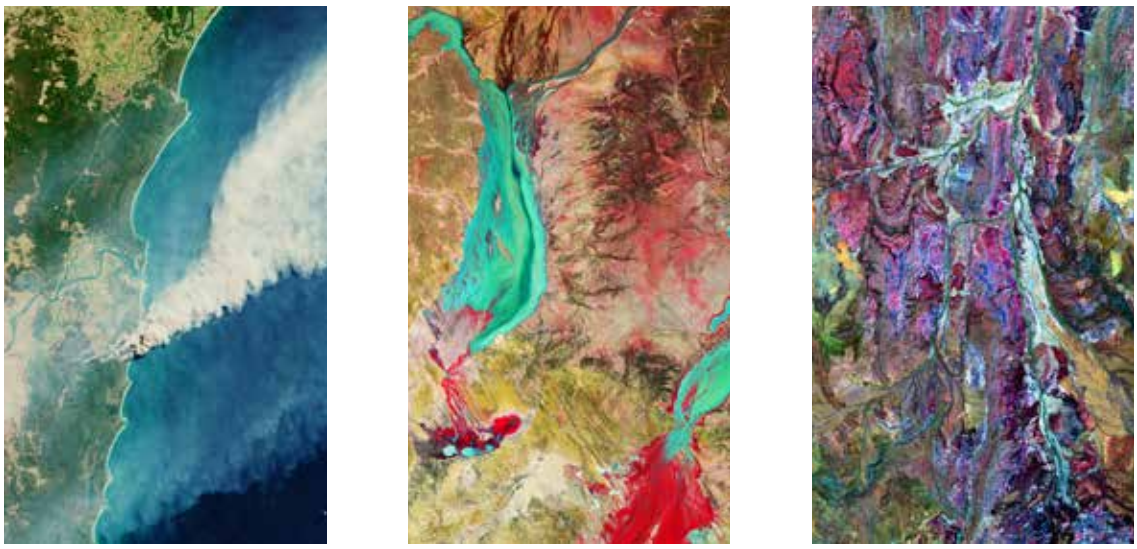

Images: Australian satellite imagery from the European Space Agency's Copernicus program stored at NCI.

We are constantly adding to our data collections' functionality to help users better access, analyse and share the data they need every day. A core aim is making sure that the data stored at NCI is safe, secure, accessible and readable. As a trusted repository for some of Australia's biggest datasets, we play a key role in the national data science community working on projects of environmental management, coastal erosion, bushfires and extreme weather.

All the significant national data collection at NCI, in total making up more than 12 petabytes of data, meet the international FAIR data standard: data should be Findable, Accessible, Interoperable and Reusable for the research community. Maintaining the datasets entrusted to us by the research community and the national science agencies for broad scientific use is a central facet of our data collections work. We are constantly working with our data providers to make sure that their data collections are up-to-date and accurate.

# SCIENTIFIC
# **VISUALISATION**

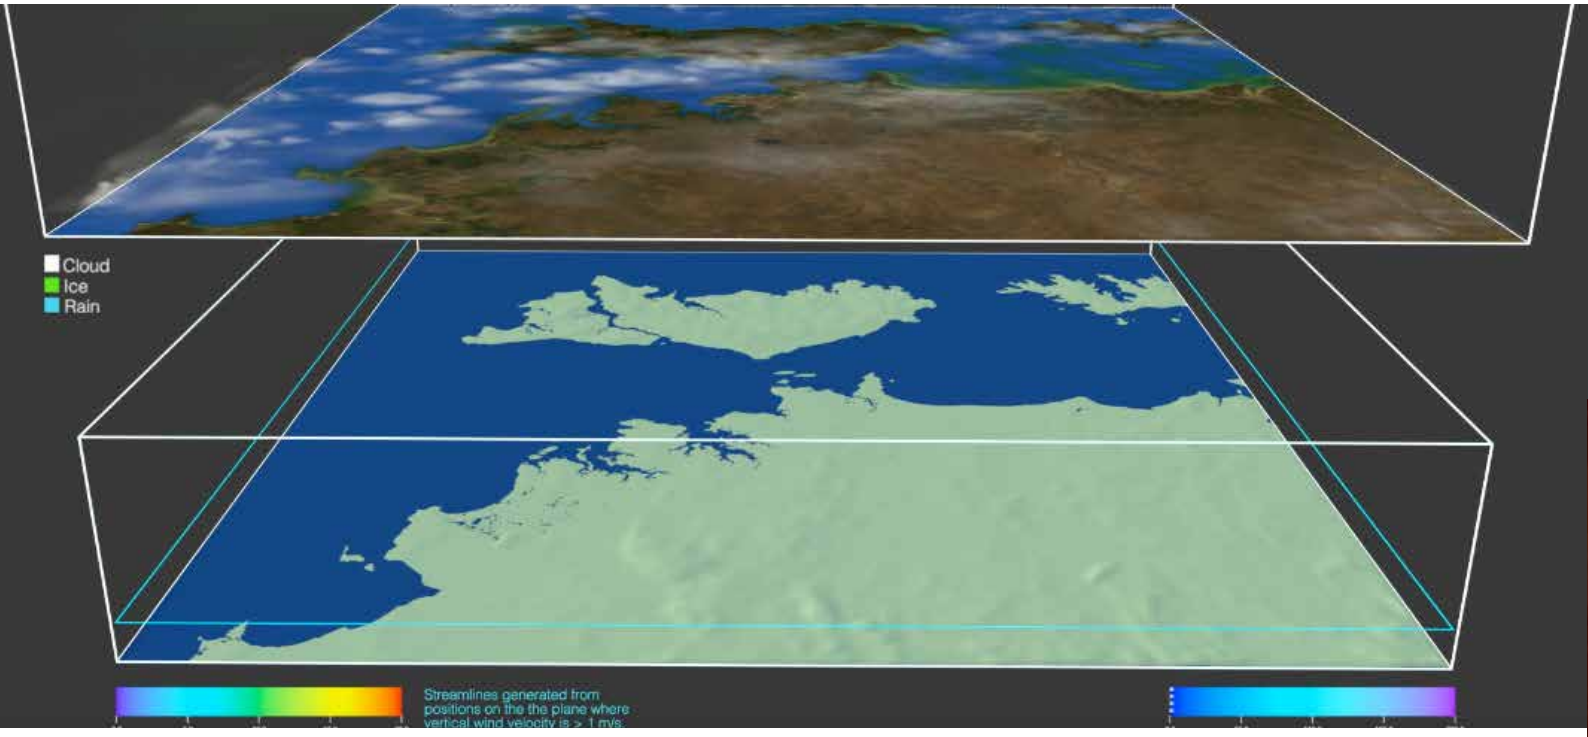


Image: Weather visualisations, VizLab NCI.

NCI's scientific visualisations are a valuable tool that enables researchers to gain deeper insights into complex datasets and easily communicate their scientific results. Produced by NCI's VizLab, these visualisations aid researchers by helping them see their data in a new way, and more easily communicate aspects of the new insights with scientific peers and the general public.

The team works closely with both researchers and in-house computational and data experts to produce scientific visualisations for the research community. These have a particular focus on weather, climate, Earth Systems science, geophysics, and environmental science, and also cover areas such as astrophysics, materials science and medicine.

By leveraging the large data storage and computational capacity available at NCI, the VizLab team works with datasets hundreds of terabytes in size to create their visualisations. Using a combination of techniques, including video animation packages used in the movie industry and NCI-developed software, NCI's scientific visualisations are an important and striking element of some of the biggest computationally-intensive research projects in Australia.

# SUPPORTING OUR USERS

NCI provides expert advice, technical support, advanced training and software installation services to the entire user community. From regular live sessions for users with pressing questions, to dedicated support for priority disciplines, NCI raises the abilities and capabilities of our users.

NCI provides a helpdesk staffed by PhD qualified computational scientists to support users with administrative and technical questions around accessing and using systems and services at NCI.

Image: NCI delivers in-person training to users.

NCI also provides an HPC-AI scholarship to PhD students across our user community, supporting early career researchers. We continue to expand on the breadth and frequency of training for our cohort of Australian researchers. We recognise the need for Australian computational scientists to develop the skills necessary to make best use of HPC and HPD infrastructure, including in growing areas of interest such as Machine Learning and bioinformatics.

We deliver regular training webinars for beginners, with an emphasis on the fundamentals of using the Gadi supercomputer. In addition, our 'Skills Sharpening' courses are an opportunity for new and established users to better grasp core HPC concepts. NCI users also have the opportunity to learn more about GPUs and their applications through bootcamps and workshops. Finally, a series of discipline-specific 'AI/ML Applications in Science' courses provide in-depth training in applying modern computational methods to scientific questions.

# COMPUTATIONAL SCIENCE **ENHANCEMENTS**

NCI's HPC Simulation, Scaling and Data Analysis Optimisation Team helps research groups understand the bottlenecks and inefficiencies in their code and plan for future growth in resolution or scale. The team can provide solutions for researchers to implement that save them time and computational resources and lead to better scientific outcomes. NCI's expertise in code optimisation starts with large climate and ocean models, and extends all the way to genomics, geophysics, fluid dynamics and astrophysics codes, as well as underlying tools beneficial to all disciplines.

For more than a decade, NCI has been supporting the Consortium for Ocean-Sea Ice Modelling in Australia (COSIMA) with the development and optimisation of some of Australia's most important scientific codes: the various forms of the Australian Community Climate and Earth-Simulation Suite (ACCESS). NCI has helped COSIMA implement improvements to the ACCESS models that have led to higher resolutions, more accurate handling of the biogeochemistry of ocean circulation and more efficient operations. The large data collections produced by and used for running software such as the global ACCESS models on tens of thousands of processors at once also requires specialised development work from NCI's team.

# ARTIFICIAL INTELLIGENCE

We are supporting Australian researchers to incorporate artificial intelligence techniques into their computational science, with dedicated training and software to enable them to solve major scientific questions. With increased data and computing capabilities provide by AI, researchers can advance their research faster than ever before.

NCI provides access to some of the latest and most powerful AI-focused hardware to researchers incorporating AI and Machine Learning techniques in their workflows. The two NVIDIA DGX A100 nodes that NCI provides are integrated into the broader Gadi supercomputer, and can interface directly with the massive amounts of data stored in our filesystems.

Image: NCI's two DGX A100 nodes from NVIDIA.

## PRECISION ONCOLOGY SCREENING PLATFORM ENABLING CLINICAL TRIALS

Alongside national medical research institutes and global pharmaceutical companies, supported by the Federal Government and enabled by supercomputing and data processing expertise, NCI is playing a central role in the Precision Oncology Screening Platform Enabling Clinical Trials (PrOSPeCT).

This groundbreaking project will use genomic methods to open up new treatment paths for people with difficult to treat cancers. PrOSPeCT will fast-track the development, manufacturing and use of precision, personalised cancer treatments.

NCI's computing power, sovereign data management and scale are critical to delivering this complex clinical, computing and data project.

In partnership with the research and industry sectors, NCI is helping make the promise of personalised medicine a reality.

# NOTES **PAGES**

# NOTES **PAGES**

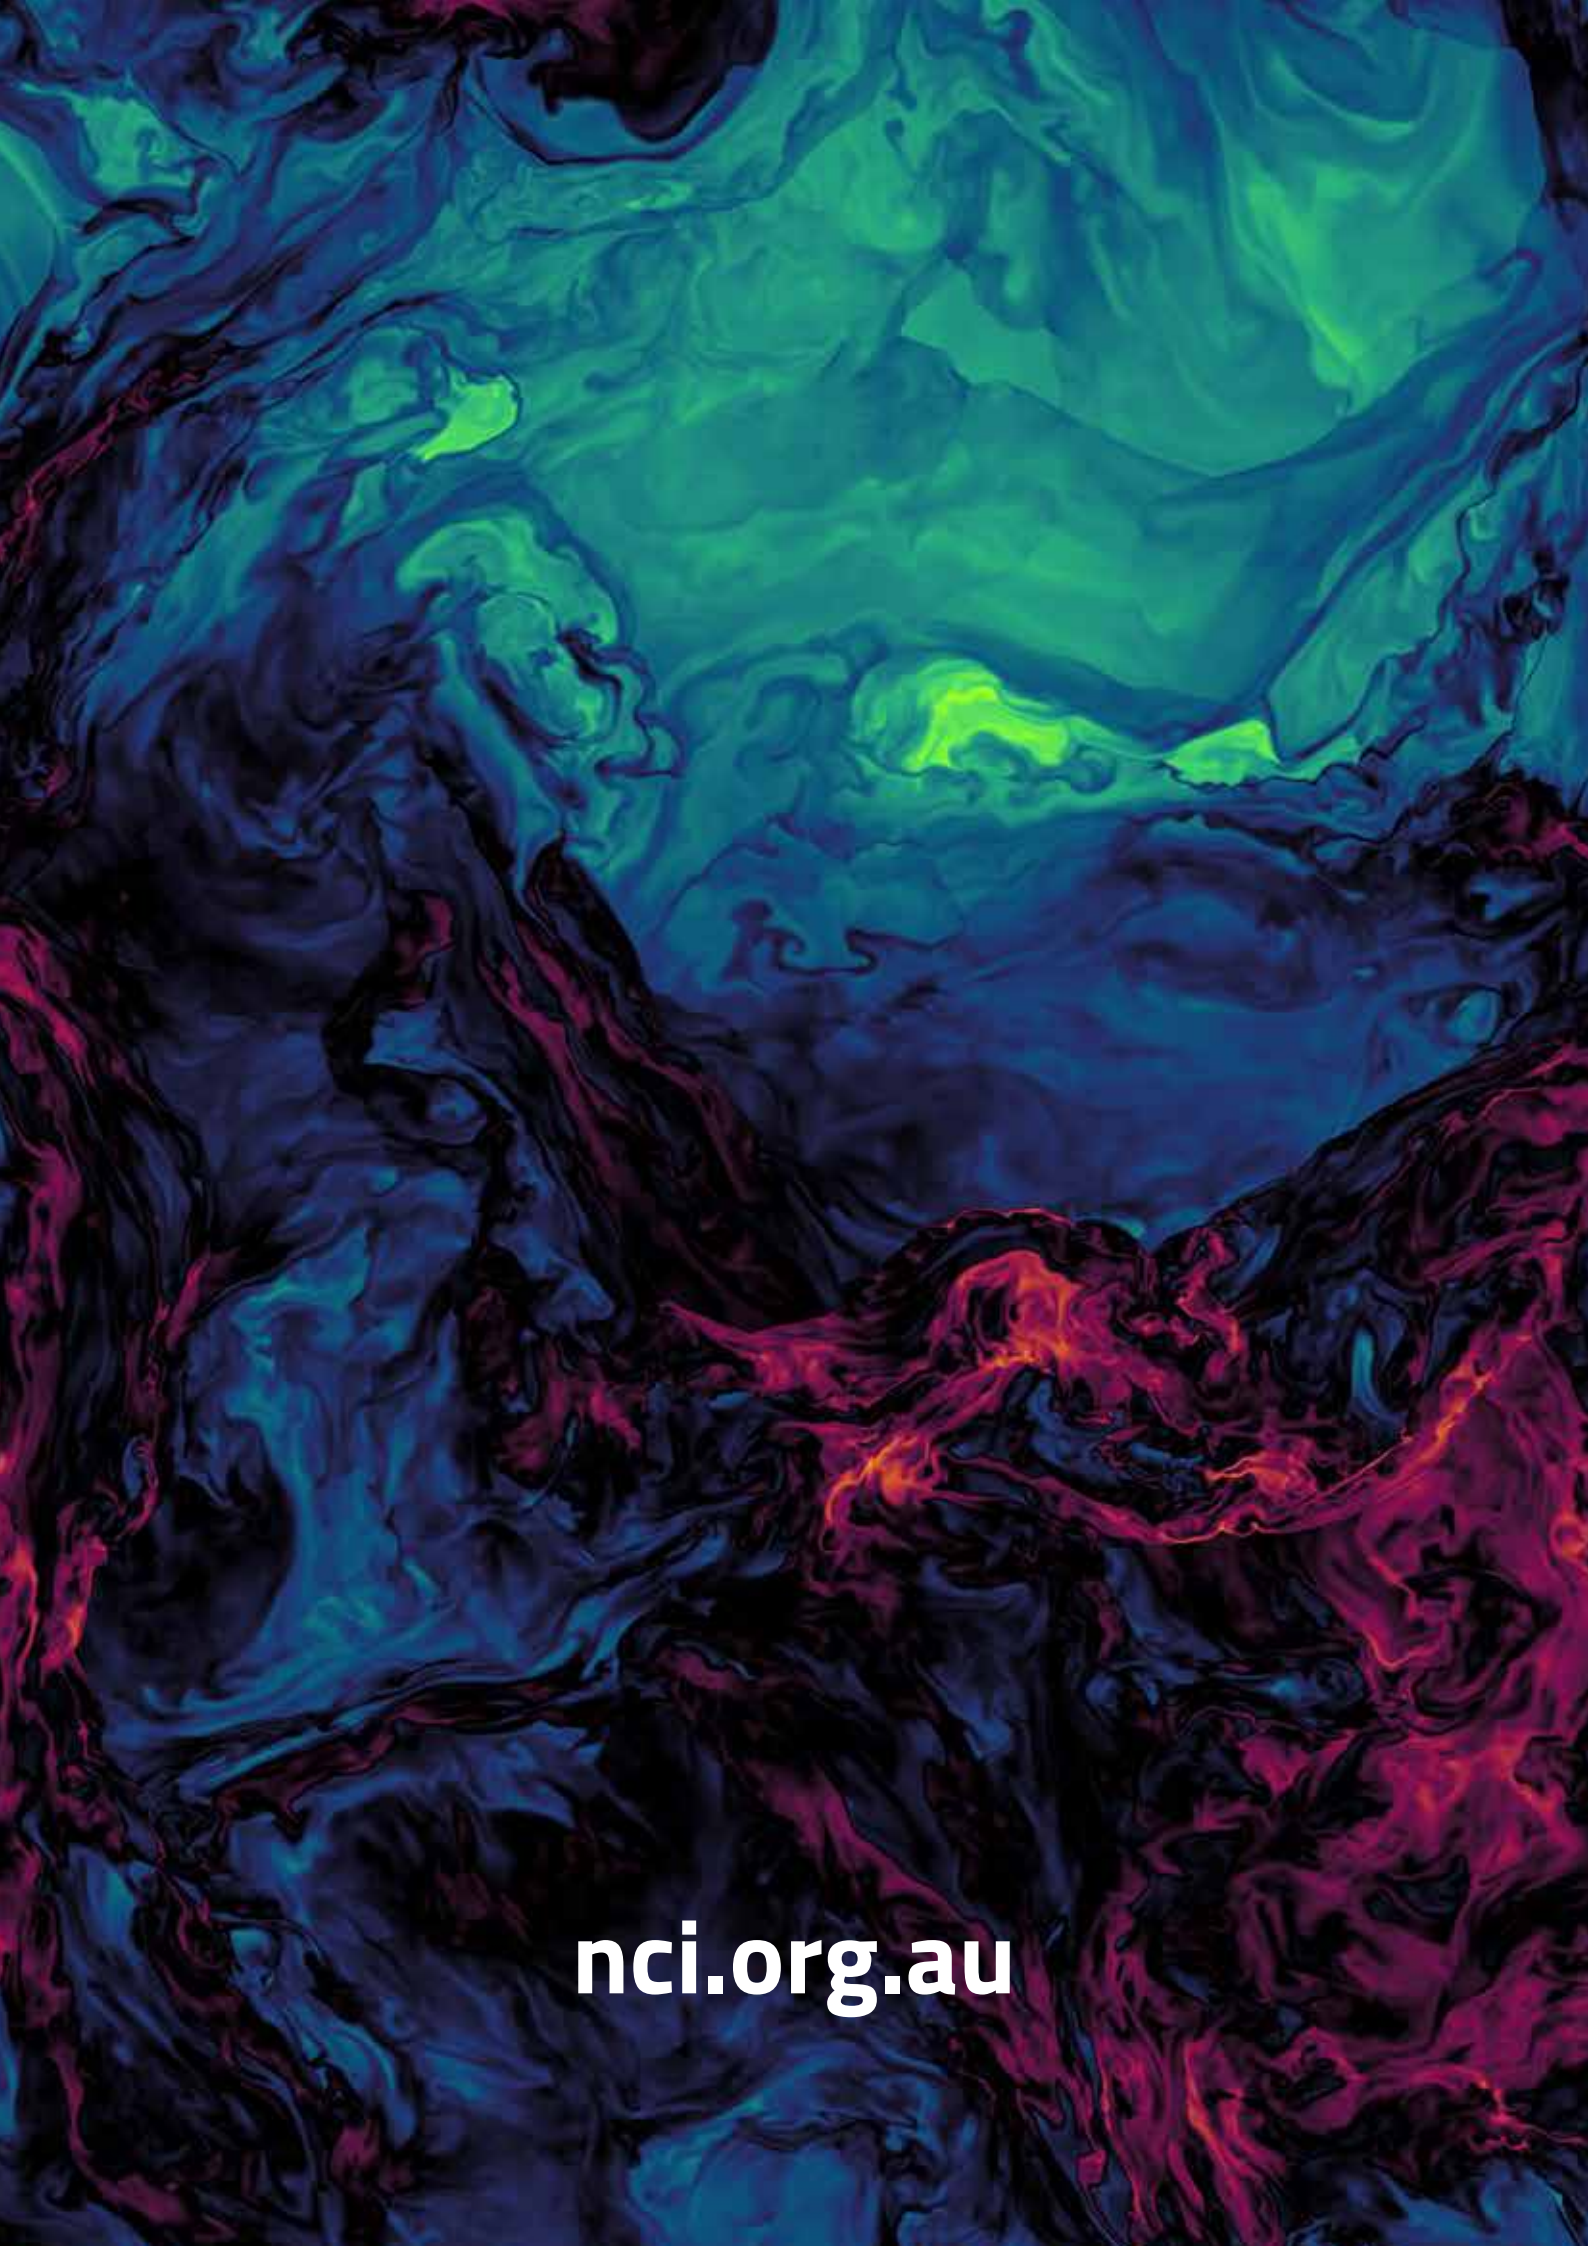

nci.org.au