

Gadi User Guide

Gadi, pronounced `gar dee`, after the words `to search for` in the language of the Ngunnawal people, the traditional owners of the Canberra region, is Australia's peak research supercomputer. The high performance computing system is hosted at National Computational Infrastructure on The Australian National University's Acton campus. Gadi was launched in Nov. 2019 and replaced its predecessor Raijin in Jan. 2020.

Gadi has 10 login nodes, 6 data mover nodes, and 184264 CPU cores in 4243 compute nodes including 160 GPU compute nodes that also provide 640 NVIDIA V100 GPUs. There is 1.67 PiB storage available locally on compute nodes and also a 22 PiB Lustre parallel filesystem in Gadi. Its inter-connect network of HDR Infiniband technology in a Dragonfly+ topology is capable of transferring data at 100 Gbps among nodes, and between nodes and the parallel filesystem. The system ran at 9264.3 TFLOP/s as of May 2020 and ranking 24th on [the June 2020 top 500 list](#).

Before starting your work on Gadi, please go through the [Welcome to Gadi](#) page and make sure you get acquainted with the basics relevant to your work.



HOW TO USE THIS USER GUIDE

The first section `Welcome to Gadi` covers most of the basics you would like to know before submitting the first job and the rest of the sections keep more in-depth topical articles and frequently asked questions in their own categories.

1. `User, Project and Scheme` focuses on the management of user account, project and schemes. It shows what can be done through [my.nci.org.au](#) and how to do it.
2. `Compute Grant and Job Debiting` helps to understand how the project compute grant is allocated and consumed by jobs.
3. `Storage and Data Management` introduces filesystems available on Gadi, and keeps our suggestions for data management and I/O operations on different filesystems.
4. `PBS Jobs` is the knowledge base for job submissions, scheduling, and monitoring.
5. `Software Applications` explains how to run software applications on Gadi
6. `Parallelism` is the showcase of jobs that benefit from the large scale parallelism provided by Gadi.

We also have two cheat sheets at part of this user guide. The `Linux Command Quick Reference` is for users who need help to get familiar with the basic commands. The `Gadi Quick Reference Guide` gathers Gadi specific facts together and can be used for quick lookups.

We are going to actively develop this user guide in the next couple of months and update it on a regular basis during the entire lifespan of Gadi. Any suggestions for its improvement are welcome. Please reach us by launching tickets at [help.nci.org.au](#) or sending an email to help@nci.org.au.