Pawsey Centre for Extreme Scale Readiness

Information for Applicants
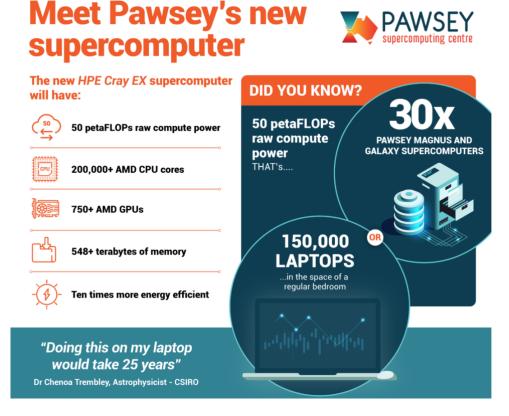
Table of Contents

Background

The next-generation Pawsey Supercomputing System (PSS) will replace Magnus and Galaxy through the [Pawsey Capital Refresh](#) project. The system will power future high-impact Australian research projects by delivering 30 times more compute power than predecessor systems Magnus and Galaxy.

Achieving sustainable scalability on PSS will create a direct pathway to achieve superior scale on next-generation supercomputers. Pawsey's Centre for Extreme Scale Readiness (PACER) is Pawsey's response to that unique opportunity.

The main purpose of PACER is to prepare Australian computational researchers for the next era of supercomputing. The focus of the PACER program is on both extreme scale research (algorithms design, code optimisation, application and workflow readiness) and using the computational infrastructure to facilitate research for producing world-class scientific outcomes.

PACER is not another allocation scheme, it is a partnership for collaboration between researchers and Pawsey Supercomputing Centre supercomputing specialists. It is built upon the past experience of Pawsey and other supercomputing facilities from running similar programs in past few years: Pawsey's [Petascale Pioneers](#) and [Athena Early Adopters](#) programs, NERSC's [NESAP](#) program and ORNL's [CAAR](#).

PACER provides researchers an opportunity to optimise their codes and workflows for next-generation supercomputers. It is designed as a minimum 3-year partnership for collaboration with Pawsey and HPC vendors by providing early access to supercomputing tools and infrastructure, training and exclusive hackathons focused on HPC performance at scale.

PACER projects will have an opportunity to create new PhD or postdoctoral positions working on technical aspects of the project in relation to new advanced architectures. The new positions can be co-funded by PACER. This co-investment model will serve as a tool to develop and grow competencies in supercomputing and computational sciences in Australia.

PACER projects will demonstrate:

- Significant performance improvement between existing Australian supercomputers and Pawsey's next-generation supercomputer,
- Significant power efficiency improvement (performance per watt),
- Superior capability improvement for the whole scientific domain and ability to achieve previously unavailable computational or data processing scales.

PACER Description

Introduction

I. **The main purpose of the Pawsey's Centre for Extreme Scale Readiness (PACER) is to prepare Australian computational researchers for the next era of supercomputing.**

The project will involve application adaptation and optimisation for selected, highly important challenges of science on a next-generation supercomputing platform available at the Pawsey Supercomputing Centre. The whole initiative is planned as a minimum 3-year collaboration project between researchers, Pawsey HPC experts, hardware and software vendors as well as other international HPC centres and experts. All projects will involve running a Grand Challenge Problem in previously unattainable scale on the next-generation supercomputer.

The focus of the PACER program is on both extreme scale research (code optimisation, application and workflow readiness) and using the computational infrastructure to facilitate research for producing scientific outcomes.

II. **Applicants are encouraged to submit high quality, innovative and original proposals**.

Proposal requirements include:

- o The computational scale of the research is exceptional,
- o The computational approach and technical project plan should cover the entire scientific discovery workflow, including supercomputing, data handling, visualisation and further processing, and storing of the scientific outcomes,
- o The results and scientific outcomes are publishable in high quality journals, and
- o The impact on the scientific domain as well as on society is significant.

III. **Applicants are encouraged to submit joint proposals involving several research groups that can benefit from the proposed computational approach**.

IV. Pawsey Supercomputing Centre will provide co-funding for postdoctoral or PhD positions (for a maximum of 2 years) embedded within a subset of successful projects to work on computational aspects of the project in close collaboration with Pawsey staff, vendors and external experts.

Applicants are required to propose a model for such a collaboration, i.e. co-funding of a postdoctoral or PhD position working part-time on the technical aspects of the PACER project in strong collaboration with Pawsey's team and vendors. This co-investment model will serve as a main tool to develop and grow competencies in HPC and computational sciences in Australia.

V. Collaboration will serve as a foundation for achieving extreme performance on next-generation supercomputers to unlock science and achieve previously unattainable scales.

PACER projects will be provided with a minimum 3-year allocation and continuous expert support from Pawsey. This will develop a long-term, close relationship and collaboration between researchers, application developers and Pawsey's supercomputing experts.

Pawsey team will become involved in software optimisation, code development as well as planning and running extreme scale computational problems. PACER program will also include the organisation of joint workshops, trainings as well as preparation of joint conference talks and journal papers. Successful applicants will be required to prepare joint research papers, participate and present at key supercomputing events and conferences in collaboration with Pawsey.

Program Stages

Detailed project schedule will be agreed case-by-case with successful applicants. However, it is assumed that all PACER projects will consist of at least the following five stages:

I.   Training

Technical and domain-specific training provided to PACER researchers and Pawsey staff . Training will discuss technical aspects of Phase-1 and Phase-2 of the PSS architectures.

II.   Preparation of the Application and Software Development Environment

Installation and testing of basic compilers, libraries and tools necessary to support all PACER projects. This stage will be implemented in close collaboration with the Vendor on the development platform and Phase-1 system. New approaches to code building and containerisation options will be evaluated at this stage as well. Pawsey staff will work closely with PACER applicants and the Vendor to define and test the application and software development environment. PACER applicants will be consulted to define their requirements, needs and confirm compatibility with their computational workflows.

III.   Code Adaptation, Porting and Benchmarking

Code adaptation and porting implemented in close collaboration between PACER researchers, Pawsey staff and other partners. PACER researchers will be required to provide reproducible test cases and unit tests together with detailed usage and result verification information. This stage will be implemented after Phase-1 becomes available and can be interpreted as a preparation to achieving superior performance on the final system delivered in Phase-2. Main dissemination activities at this stage will include workshops, conference talks and technical papers.

IV.   Code Optimisation

Code optimisation implemented in close collaboration between PACER researchers, Pawsey staff and other partners. This stage will be implemented after Phase-2 becomes available in the preparation of Grand Challenge Problems. The significant performance improvement to existing Australian supercomputers will be demonstrated at this stage. Main dissemination activities at this stage will include workshops, conference talks, publications in computational science journals as well as multidisciplinary publications in high profile journals.

V.   Grand Challenge Problems

Grand Challenge Problems will be run on the Phase-2 system to demonstrate superior capability improvement for the scientific domain as well as ability to achieve previously unavailable computational and processing scales. Main dissemination activities at this stage will include conference talks as well as publications in computational science, domain-specific as well as multidisciplinary, high profile journals.

Collaboration and Co-funding

The main purpose of PACER is to develop a long-term, close partnership and collaboration between research groups, application developers and Pawsey's supercomputing experts. PACER projects are expected to:

- prepare and present joint conference papers and talks,
- prepare and publish joint multidisciplinary publications in high profile journals,
- prepare and publish joint technical papers,
- present PACER results on major supercomputing conferences and events,
- co-organise workshops and regular project meetings.

It is also expected that Pawsey staff members assigned to the project will co-author publications and conference talks.

To achieve these goals PACER projects might apply for co-funding of a PhD or postdoctoral position working on both research and technical parts of the project (for a maximum of 2 years). We encourage the applicants to propose a specific model of such co-funding since we acknowledge that this might differ for each institution. The preferred model is the PhD or postdoctoral position hired at the applicants' institution with a MoU agreement defining the scope and evaluation periods of the work to be performed.

Important Dates

Call Opens	26th October 2020
Call Closes	29th November 2020
Proposals Assessment	29th November 2020 – 3rd January 2021
Shortlisting and Program Committee Meetings *	3rd January – 22nd January 2021
Applicants Notified	End of January 2021

* At this stage shortlisted Applicants might be contacted and asked to provide additional information e.g. in relation to co-funding and collaboration proposal.

Eligibility Criteria

Projects must be research-focused.

PACER is not an allocation scheme, the proposals should be focused on achieving previously unavailable scale of simulations or data processing workflows in the given scientific domain rather than running regular research computations.

Applicants are expected to have demonstrated experience **using supercomputing resources at large scale** at an Australian or international facility. Proposals have to demonstrate significant computational or data processing challenges that need access to supercomputers at scale.

Applicants are required to propose a co-investment and collaboration model as well as demonstrate alignment with Australian Science and Research Priorities and Programs.

Principal Investigators (PIs) of projects must hold at least a 0.2 FTE research position at an Australian higher-education institution, research institute, publicly funded research agency or Australian Government agency conducting research & development, and be based in Australia.

A person holding a Postdoctoral appointment is eligible to be a Principal Investigator on a PACER application.

An individual may be named as a Principal Investigator on only one PACER application.

Application Procedure

Submission forms are available in MS Word and LaTeX formats at PACER program webpage:
http://pawsey.org.au/pacer

Completed forms should be submitted in PDF format by emailing: pacer@pawsey.org.au.

Instructions on what information and content need to be provided is given in the submission form for each of the submission part.

Review Process

PACER Committee

PACER Committee will be responsible for conducting a fair review and selection process based on the selection criteria.

PACER Committee is composed of:

- PACER Program Committee,
- Scientific Review Panel,
- Technical Review Panel,
- Collaboration Review Panel.

All above reviewing entities are composed of researchers, supercomputing specialists and professionals from Pawsey Supercomputing Centre, Australian Research and Government Organisations as well as international Supercomputing Centres and Research Organisations.

The expertise of the Scientific Review Panel covers a broad range of scientific disciplines including (but not limited to): engineering, computational chemistry and physics, radio astronomy, and life sciences.

Assessment Criteria and Scoring

The selection of projects will be based on the following assessment criteria:

Research Review - 40%

- Research significance - 20%
- Research track record - 20%

Collaboration and Co-Investment Review – 10%

- Collaboration, co-investment, capability and sufficient effort – 10%

Science and Research Priorities and Programs Review – 10%

- Alignment with National Science and Research Priorities and Programs – 10%

Technical Review - 40%

- Technical capability - 10%
- Technical track record - 10%
- Capabilities of codes - 10%
- Technical project plan - 10%

Portfolio Approach

The Program Committee, when assessing PACER proposals will be taking a portfolio approach to support a wide spectrum of scientific disciplines, especially those covered by National Science and Research Priorities and Programs (https://www.industry.gov.au/data-and-publications/science-and-research-priorities). The Program Committee can request two (or more) projects to combine.

Appeal Process

All decisions of the PACER Program Committee are final. Appeals will be considered only against administrative or procedural issues and not against decisions of the Program Committee or against assessor ratings and comments.

Allocations on Pawsey Systems

PACER is not another allocation scheme, it is a partnership for collaboration between researchers and Pawsey Supercomputing Centre supercomputing specialists. Therefore, all applicants are encouraged to apply for separate NCMAS, Pawsey Partner or Energy and Resources allocations within available calls for projects.

All PACER projects will be provided with minimum three-year allocation on current and next-generation Pawsey supercomputing systems strictly for the following purposes:

- Training and acquiring new supercomputing skills on next-generation systems,
- Development and optimisations of codes and workflows for the need of PACER projects,

- Running reference simulations or data processing workflow on supercomputing resources currently available at Pawsey Supercomputing Centre (Magnus, Galaxy, Zeus or Topaz),
- Running large scale profiling jobs for the need of PACER projects,
- Running Grand Challenge Problem on the next-generation supercomputer at Pawsey.

The allocation size will be flexible to achieve goals defined in PACER project proposals. An allocation of maximum 1M Service Units (SUs) per year will be granted for PACER projects with additional dedicated Early Access to Phase 2 system for the purpose of running Grand Challenge Problems.

Getting Help

For assistance with your application or for any other enquiries, please contact the Pawsey help desk via the User Support Portal or via email at help@pawsey.org.au.