# OOKAMI PROJECT APPLICATION

**Date:** **24/04/2021**

**Project Title: MIT General Circulation Model Performance Tuning on ARM64 v8 SVE**

**Usage:**

☒ Testbed

☐ Production

**Principal Investigator: Christopher Hill**

University/Company/Institute: Massachusetts Institute of Technology

Mailing address including country: 54-1524 MIT, 77 Mass. Ave, Cambridge, MA02139, USA

Phone number: 6172536430

Email: cnh@mit.edu

**Names & Email of initial project users:**

Chris Hill, cnh@mit.edu

**Usage Description:**

We wish to examine the performance of a benchmark calculation ( https://youtu.be/CCmTY0PKGDs ) of the MIT General Circulation Model ( https://mitgcm.org) on an ARM v8 + HBM2 system. The MITgcm code ( https://github.com/mitgcm/mitgcm ) and algorithm is a widely used open-source geophysical fluids tool. It is used in between 20-30 publications every month by researchers world-wide. It is used in everything from single core processes to large MPI simulations running on tens of thousands of cores. We are interested in evaluating how well the code performs on an A64 FX cluster system and updating the MITgcm code repository to provide a recipe for A64 FX execution.

**Computational Resources:**

Total node hours per year: 5,000 node hours/year max

Size (nodes) and duration (hours) for a typical batch job: Most runs will be short (< 1 hour). Between 1 and 16 nodes will be used.

Disk space (home, project, scratch): 2TB

**Personnel Resources:**

**Required software:**
C and Fortran compilers, MPI libraries, NetCDF or HDF if possible.

**If your research is supported by US federal agencies:**

Agency: NSF, NASA, DOD.

Grant number(s):
NSF- **1640831**

NSF- **1835778**

DOD-AirForce - #FA8750-19-2-1000

NASA - CREI 1572041

---

**Production projects:**

Production projects should provide an additional 1-2 pages of documentation about how
(a) the code has been tuned to perform well on A64FX (ideally including benchmark data comparing performance with other architectures such as x86 or GPUs)
(b) it can make effective use of the key A64FX architectural features (notably SVE, the high-bandwidth memory, and NUMA characteristics)
(c) it can accomplish the scientific objectives within the available 32 Gbyte memory per node