

OOKAMI PROJECT APPLICATION

Date: 4/13/23

Project Title: Evaluation of Navy-Relevant codes on ARM cluster

Usage:

Testbed

Production

Principal Investigator: Robert Rosenberg
University/Company/Institute: Naval Research Laboratory
Mailing address including country: 4555 Overlook Ave, Washington, DC 20375
Phone number: (202) 404-7495
Email: robert.rosenberg@nrl.navy.mil

Names & Email of initial project users:

Robert Rosenberg	robert.rosenberg@nrl.navy.mil
Keith Obenschein	keith.obenschain@nrl.navy.mil
Yu Yu Khine	yuyu.khine@nrl.navy.mil
Owen Lagarde	owen.lagarde.ctr@nrl.navy.mil
Gopal Patanaik	gopal.patnaik.ctr@nrl.navy.mil

Usage Description:

The goal of this project is to determine if Navy-relevant codes perform better with high-bandwidth memory. The focus will be on codes with low arithmetic intensity, which are expected to perform better with high memory bandwidth. We will be benchmarking various test codes and codes of Navy interest in order determine what combination of compilers, compiler switches and code optimizations work best. We will compare performance against conventional clusters to determine if there is a benefit to high-bandwidth memory. Only open source codes will be tested.

Computational Resources:

Total node hours per year: 500

Size (nodes) and duration (hours) for a typical batch job:
4 nodes 0.25 hours
(1-32 nodes for scaling studies)

Disk space (home, project, scratch):

10GB,0GB,120GB

Personnel Resources (assistance in porting/tuning, or training for your users):

Assistance in tuning.

Required software:

Compilers : gcc, nvhpc, fujitsu, armflang

Editors : vim, emacs

Parallel Environments : MPI, OpenMP

If your research is supported by US federal agencies:

Agency: DoD HPC Modernization Program

Grant number(s): N/A

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