OOKAMI PROJECT APPLICATION

Date: 8/3/2021

Project Title: Testing/Tuning WRFv3.9.1.1 on A64FX

Usage:

☒ Testbed

☐ Production

Principal Investigator: Evan Dodge

University/Company/Institute: NASA Advanced Supercomputing Division, NASA Ames Research Center

Mailing address including country: NASA Advanced Supercomputing Division, NASA Ames Research Center, Mailstop 258-6, Moffet Field, CA 94035-1000, USA

Phone number: 650-604-4356

Email: evan.k.dodge@nasa.gov

Names & Email of initial project users: Evan Dodge: evan.k.dodge@nasa.gov , Johnny Chang: johnny.chang@nasa.gov

Usage Description: We will be compiling, running, and tuning WRF v3.9.1.1 on the A64FX platform in order to compare with Intel Xeon results on the Xeon Gold 6148 and 6248 processors. We will be using the Typhoon Morakot dataset which NAS uses to benchmark systems at NASA Ames Research Center.

Our colleague, Samson Cheung, started this project in late June, 2021 on Ookami. He is no longer continuing with the project and we will be taking over. We are also requesting access to his files.

Computational Resources:

Total node hours per year: 3,000 (60 node hrs/run * 50 runs includes tuning)

Size (nodes) and duration (hours) for a typical batch job: 30 nodes, 2 hrs

Disk space (home, project, scratch): We do not anticipate needing more than 1TB total

Personnel Resources (assistance in porting/tuning, or training for your users**): None**

Required software: C++/Fortran compilers, MPI

If your research is supported by US federal agencies:

    Agency: NASA Ames Research Center

    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