OOKAMI PROJECT APPLICATION

Date: September 17, 2021

Project Title: Chapel on Ookami

Usage:

☒ Testbed

☐ Production

Principal Investigator: Bradford Chamberlain

University/Company/Institute: Hewlett Packard Enterprise

Mailing address including country:

Hewlett Packard Enterprise

901 Fifth Ave, Suite 1100

Seattle, WA 98164 USA

Phone number: 206-701-2077

Email: bradcray – at – hpe – dot – com

Names & Email of initial project users:

- Elliot Ronaghan <elliot.ronaghan – at – hpe – dot – com>
- Engin Kayraklioglu <engin – at – hpe – dot – com>

Usage Description:

Chapel is a parallel programming language designed with a modern sensibility to make scalable parallel programming far more productive than with current broadly adopted approaches—think "as programmable as Python yet as scalable as C++/Fortran + MPI + OpenMP." Our intention is to use our time on Ookami to ensure that Chapel is portable to HPC systems that utilize A64FX processors, to do performance analysis and tuning experiments at larger scales than we can do in-house, and to develop best practices for Chapel programmers on Ookami and similar systems. We will also use the account to help support and collaborate with any other users who may use Chapel on Ookami going forward.

Computational Resources:

Total node hours per year: ~1000 max?

Size (nodes) and duration (hours) for a typical batch job: The vast majority of our runs tend to focus on a modest numbers of nodes (1–16) followed by scalability sweeps on powers of two up to the max available (so, 1, 2, 4, …, 128, 174). Most of our benchmarking and debugging runs are reasonably short (small numbers of minutes). If

we were to try some user codes and workloads, those can be longer-running depending on I/O rates—perhaps 30-60 minutes. But these are relatively rare.

Disk space (home, project, scratch): 40 GB, 4TB, 4TB

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

Required software: Nothing special: C/C++ compilers, Pthreads, job launchers, …

If your research is supported by US federal agencies:

Agency: N/A

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