

OOKAMI PROJECT APPLICATION

Date: June 21, 2022

Project Title: Parfu Parallel Storage Tool Benchmarking and Testing

Usage:

- Testbed

Principal Investigator: Craig P Steffen

- University/Company/Institute: University of Illinois
- Mailing address including country:
Craig Steffen
1205 West Clark St
Urbana IL 61821
- Phone number: 217-300-1881
- Email: csteffen@gmail.com

Names & Email of initial project users:

Craig Steffen csteffen@illinois.edu

Usage Description:

I will be compiling the 1.0 version of the test code, parfu, and dtar, and any other direct competitors I find. I will then benchmark each code on each of the three test collections for different configurations on one node to find the optimum configuration. Then further jobs will be run to test that configuration at larger scale until they no longer scale well; this may go up to 8 nodes or perhaps 16 nodes, but likely not past that. Runs tend to run 20 minutes or so at the longest, maybe up to an hour or so on very slow configurations/codes.

Computational Resources:

- Total node hours per year: 500
- Size (nodes) and duration (hours) for a typical batch job: Lots of single-node jobs, then scaling up to perhaps 8 or maybe 16 nodes.
- Disk space (home, project, scratch): Test data (3 collections) will fit in 2 TB (it's over 1.5 million files). Each job will have probably 5 archive files of up to 1 TB each, but they can be on /scratch and deleted immediately after the job.

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

None.

Required software:

An MPI C++ compiler and corresponding MPI run-time.

If your research is supported by US federal agencies:

- Agency: NSF
 - Grant number(s): 1238993
-