

FPROF: Parallel Namespace Tree Walk Benchmark

Summary Version

1.0

Purpose of Benchmark

fprof is a lightweight profiler designed to provide a set of useful statistical characterization of the target file system at extreme scale. fprof can be used for benchmarking the parallel namespace walk efficiency and performance.

Characteristics of Benchmark

fprof uses Python and MPI.

Location of Benchmark

<https://github.com/olcf/pcircle>

Mechanics of Building Benchmark

see README.md at above repo.

Mechanics of Running Benchmark

```
mpirun -np 8 fprof /path/to/fs
```

Further run instruction and sample output are given in the README from repo.

Verification of Results

The FOMs reported are (1) estimated or actual time required to walk a global namespace with 1 billion directories and 1 billion files on a current system, and (2) estimated time required to walk a global namespace with 5 billion directories and 5 billion files on the proposed system (both in seconds).