CAM-SE

Summary Version
1.1

Purpose of Benchmark
This goal of this benchmark is to determine the throughput rate (Simulated Years Per Day, SYPD) and parallel efficiency using a large portion of the proposed machine. The benchmark is developed to run on Mira.

Characteristics of Benchmark
The Community Atmosphere Model – Spectral Element (CAM-SE) is an actively developed atmospheric climate modeling code that is widely used by climate scientists as the default atmospheric model of the Community Earth System Model (CESM), used also for climate projections in the Intergovernmental Panel on Climate Change (IPCC). The "dynamical core" of CAM-SE is called HOMME (High-Order Methods Modeling Environment), and this will be used for the RFP. This models the stratified, compressible, hydrostatic Euler equations on the sphere with added multi-scale physics representing climate-related processes. It is discretized with a Spectral Element method on the cubed-sphere grid. Inside CESM, it is coupled with other model components. It has been hybrid MPI and OpenMP parallelized and efficiently scales to more than half of Titan. The benchmark uses a 14km grid, representing at-scale scientific capability at very high resolution.

Mechanics of Building Benchmark
Use the `compile_and_run.mira` script located in the `homme1_3_6` folder, and his will run all of the benchmarks on Mira by compiling the code and then submitting jobs from the `test/ASP/baroclinic.mira` folder.

Mechanics of Running Benchmark
This problem will be run at three node counts in a strong scaling study: 1024 nodes, 2048 nodes, and 4096 nodes. Perfect scaling would mean the SYPD throughput would reduce by the same factor as the node count increase. Also, parallel efficiency and MPI waiting time will be diagnosed. The scaling figure of merit is reported in SYPD.

Verification of Results
The figure of merit (FOM) is simulated years per (wallclock) day (SYPD). The FOM can be generated from the timing files using the postproc_scaling.sh script. This script should be machine independent after installing gnuplot and changing the directories of the timing files.

Correct answers for the single-node problem size can be verified by a diff of the output against a reference output, to confirm machine precision differences. As the

tracers do not interact non-linearly in the benchmark, error growth should be attributed to successive round-off differences. The file for comparison is asp_baroclinic2_1node_compare.nc.