



Addressing capability computing challenges of high-resolution global climate modelling at the Oak Ridge Leadership Computing Facility

Valentine Anantharaj (1), Matthew Norman (1), Katherine Evans (2), Mark Taylor (3), Patrick Worley (2), James Hack (1), and Benjamin Mayer (2)

(1) National Center for Computational Sciences, Oak Ridge National Laboratory, Oak Ridge, United States, (2) Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, United States, (3) Sandia National Laboratories, Albuquerque, United States, (4) Climate Change Science Institute, Oak Ridge National Laboratory, Oak Ridge, United States

During 2013, high-resolution climate model simulations accounted for over 100 million “core hours” using Titan at the Oak Ridge Leadership Computing Facility (OLCF). The suite of climate modeling experiments, primarily using the Community Earth System Model (CESM) at nearly 0.25 degree horizontal resolution, generated over a petabyte of data and nearly 100,000 files, ranging in sizes from 20 MB to over 100 GB. Effective utilization of leadership class resources requires careful planning and preparation. The application software, such as CESM, need to be ported, optimized and benchmarked for the target platform in order to meet the computational readiness requirements. The model configuration needs to be “tuned and balanced” for the experiments. This can be a complicated and resource intensive process, especially for high-resolution configurations using complex physics. The volume of I/O also increases with resolution; and new strategies may be required to manage I/O especially for large checkpoint and restart files that may require more frequent output for resiliency. It is also essential to monitor the application performance during the course of the simulation exercises. Finally, the large volume of data needs to be analyzed to derive the scientific results; and appropriate data and information delivered to the stakeholders.

Titan is currently the largest supercomputer available for open science. The computational resources, in terms of “titan core hours” are allocated primarily via the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) and ASCR Leadership Computing Challenge (ALCC) programs, both sponsored by the U.S. Department of Energy (DOE) Office of Science. Titan is a Cray XK7 system, capable of a theoretical peak performance of over 27 PFlop/s, consists of 18,688 compute nodes, with a NVIDIA Kepler K20 GPU and a 16-core AMD Opteron CPU in every node, for a total of 299,008 Opteron cores and 18,688 GPUs offering a cumulative 560,640 equivalent cores. Scientific applications, such as CESM, are also required to demonstrate a “computational readiness capability” to efficiently scale across and utilize 20% of the entire system. The 0,25 deg configuration of the spectral element dynamical core of the Community Atmosphere Model (CAM-SE), the atmospheric component of CESM, has been demonstrated to scale efficiently across more than 5,000 nodes (80,000 CPU cores) on Titan. The tracer transport routines of CAM-SE have also been ported to take advantage of the hybrid many-core architecture of Titan using GPUs [see EGU2014-4233], yielding over 2X speedup when transporting over 100 tracers. The high throughput I/O in CESM, based on the Parallel IO Library (PIO), is being further augmented to support even higher resolutions and enhance resiliency. The application performance of the individual runs are archived in a database and routinely analyzed to identify and rectify performance degradation during the course of the experiments.

The various resources available at the OLCF now support a scientific workflow to facilitate high-resolution climate modelling. A high-speed center-wide parallel file system, called ATLAS, capable of 1 TB/s, is available on Titan as well as on the clusters used for analysis (Rhea) and visualization (Lens/EVEREST). Long-term archive is facilitated by the HPSS storage system. The Earth System Grid (ESG), featuring search & discovery, is also used to deliver data. The end-to-end workflow allows OLCF users to efficiently share data and publish results in a timely manner.