CESM

Community Earth System Model

Proposal for CSL Resources

PI: Jean-François Lamarque, CESM Chief Scientist
Period of Performance: 11/1/16 – 10/31/18
Total Request: 250M Yellowstone Core Hours
420M Cheyenne Core Hours

Table of Contents

Introduction

The Community Earth System Model (CESM) project is a community effort that requires collaboration between scientists from universities, national laboratories, and other research organizations to continuously develop, test, improve, and apply a comprehensive Earth modeling system. In recent years, this process has been almost exclusively facilitated through access to Climate Simulation Laboratory (CSL) computational resources. The CESM and its predecessor, the Community Climate System Model (CCSM), have been at the forefront of international efforts to understand and predict the behavior of Earth's climate. Evaluation of the CESM1-CAM5 has ranked it amongst the best climate models in the world (Knutti et al., 2013). Output from numerous simulations using CCSM and CESM are routinely used in many hundreds of peer-reviewed studies to better understand the processes and mechanisms responsible for climate variability and change. Significant CSL-supported efforts such as the CESM Large Ensemble have been key in advancing our understanding of the climate system. CESM source code and simulation output are made freely available to the broad scientific community. Additionally, CCSM and CESM simulations have generated important contributions to both national and international assessments of climate, including those of the Intergovernmental Panel on Climate Change (IPCC) and the U.S. Global Change Research Program (USGCRP). CESM provides NSF and DOE, its primary sponsors and partners in the USGCRP, and the national and international research communities with a well-supported core modeling system for multiple purposes, including studies of past and current climate, and projections of future climate change.

With the advancement brought about through the development of CESM2, community involvement in CESM development and application has continued to expand. Accordingly, the objectives and priorities outlined in this proposal emanate directly from the community of scientists who participate in the CESM project – the 12 CESM working groups and the CESM Scientific Steering Committee (SSC, whose membership consists of not only NCAR scientists but also scientists from universities and government laboratories). In particular, to prepare this proposal, each working group consults with their constituents (beginning at the June 2016 Breckenridge CESM Workshop, and with widely distributed emails) to discuss model development goals and production simulations required to address high priority scientific questions, especially those that benefit from analysis and interpretation by the broader community. This resulted in draft working group plans. This collection of draft plans was then distributed and reviewed, revised, refined and prioritized through a process of exchange across the different working groups, with the goal of producing a coherent and coordinated plan for the use of the CSL resource over the upcoming period of performance. The plans and resource requests of the individual working groups and community projects, which appear as appendices, then served as the source material for further deliberation by the CESM SSC. The goal of the SSC in this proposal was to articulate the overarching development and production simulation priorities for the entire CESM project, as well as a more detailed description of the main development and production activities, the required computing resources, an estimate of the amount of data to be generated, and a management plan to deal with the data volume. During the previous CESM proposal preparation, a similar process was implemented and we believe resulted in a coherent overview of the testing, development and application needs of the broad CESM project.

Overarching Priorities

Over the period of performance for the present request, the main priorities for the utilization of CSL resources will be

1. Release of CESM2 (scheduled for early 2017); this version contains many significant improvements in all components, including a land-ice model.
2. Perform base CMIP6 (Coupled Model Intercomparison Project, Phase 6) simulations on Yellowstone in 2017 (mostly using the CESM2 1^{o} version)
3. Development and testing of the CESM2 high-resolution configurations
4. Development of new ocean model (prompted by changes to longstanding institutional arrangements and external programmatic decisions).

While the release of CESM2 was previously scheduled for July 2016, it was decided by the CESM SSC (in agreement with the working groups) that recent model developments, such as the introduction of a new representation of cloud physics in 2015 and major changes throughout the land model component of CESM, needed more time to provide significant improvements over the existing CESM1 configurations. Such improvements have now been achieved (see the AMWG section in the Accomplishment document, for example), and we are confident that the new simulations will demonstrate a significant advancement in model skills as well as capability for the broad CESM community. In following internationally established CMIP protocols, specific simulations (e.g. a long pre-industrial control and at least 1850-present coupled simulation) are required for the scientific release of CESM2. Note that these simulations are also part of the core simulations for CMIP6 (referred to as the DECK simulations, see below) and this proposal therefore represents an efficient use of computational resources by serving multiple purposes. Furthermore, coordination with CMIP6 will elevate the visibility of CESM, as described next.

The CMIP6 simulations consist of 1) the DECK simulations (long pre-industrial control, 1%-increase in CO_2, instantaneous increase to $4xCO_2$, and a specified sea-surface temperature simulation, 1979-present) and 2) some of the proposed simulations (focusing on highest priority Tier 1 experiments) from several different Model Intercomparison Projects (MIPs), each of which has been carefully reviewed by the working groups with the approval of the SSC. Note that the DECK simulations have to be performed for each configuration that will be part of CMIP6 (see discussion below). All CMIP6 simulations with the 1° version of CESM2 will be performed on Yellowstone (focusing on the period 1/1/17-12/31/17). *By participating in CMIP6 (in the DECK and MIPs), the 1° released version of CESM2 will be subject to a broad and intense scrutiny by the national and international scientific communities, thereby providing a high level of documentation on a model that the CESM Community will be able to use for many years.* In addition, it will provide data on a wide range of scientific areas for climate change and related impacts research.

Beyond the 1° version of CESM2, experience will be gained from performing simulations at finer horizontal and vertical resolutions for the atmosphere, ocean, sea-ice and land surface. One of the main drivers for pursuing high-resolution (globally or with regional mesh refinement) simulations is the ability to improve the representation of mesoscale processes (such as oceanic eddies, tropical cyclones, atmospheric rivers, topographically-forced circulations, …). In particular, CAM5 simulations found that midlatitude winter and spring extreme precipitation over land are significantly better represented in the high-resolution model configuration (Wehner et al., 2104). Also, the combination of high-resolution in the ocean and atmosphere enables the representation of important small-scale features such as air-sea interaction over ocean frontal zones (Small et al., 2014). At the same time, it must be recognized that global horizontal resolution (25-km grid spacing or finer) simulations pose significant challenges due to high computational cost and the poor performance of some of the physical parameterizations when translated from low-resolution configurations (Bacmeister et al., 2104).

In addition, it has been recognized that the current ocean component of CESM (POP-2) will not be developed further under the Department of Energy's earth-system science activities; this has prompted a review of the future of ocean model development within and beyond the CESM Community. The SSC has instructed an independent panel to review various options for ocean modelling that will be required for future versions of CESM, starting with CESM3 (see the OMWG discussion in the Supplemental Material for more details). This discussion has immediate implications for computational resources related to ocean model development, and will further affect sea-ice and land-ice modeling (see OMWG and PCWG proposals).

CMIP6 and its Relation to CESM2 Release

The CMIP6 experimental design can be found at http://www.wcrp-climate.org/wgcm-cmip/wgcm-cmip6 and is discussed in more details in Eyring et al. (2016). With the Grand Challenges of the World Climate Research Programme (WCRP) as its scientific backdrop, CMIP6 will focus on three broad questions:

– How does the Earth system respond to forcing?
– What are the origins and consequences of systematic model biases?

– How can we assess future climate changes given internal climate variability, predictability, and uncertainties in scenarios?

In order to address those questions, CMIP6 consists of two major elements: (1) the DECK (Diagnostic, Evaluation and Characterization of Klima) and CMIP historical simulations (1850–near present) that will maintain continuity and help document basic characteristics of models across different phases of CMIP; and (2) an ensemble of CMIP-Endorsed Model Intercomparison Projects (MIPs) that builds on the DECK and CMIP historical simulations. These MIPs target a range of specific questions and aim to fill the scientific gaps of the previous CMIP phases. Participation in CMIP6-Endorsed MIPs is a matter of scientific judgement and therefore at the discretion of the participants in each modeling group. Scientific prioritization in the CESM project takes place via its working-group structure with the overall approval of the CESM SSC. This process has led to the listed selection of CMIP6 MIPs (see Table 1) that are deemed of sufficient scientific interest to be performed using the 1° version of CESM2 primarily. Under the specifically labeled CMIP6 portion of the current proposal (see below), only the DECK and Tier 1 (highest priority) simulations are covered, except for a 9-member ensemble of 21st century projections (Tier 2). Any other Tier 2 and Tier 3 simulation that working group members are interested in performing are included in their respective requests.

The complete collection of CESM simulations (DECK + Tier 1) is available as https://www.dropbox.com/s/a1zdy4ej1l4n842/deck_tier1_endorsed_20160901.xlsx?dl=0. From this list, it is clear that the participation of CESM to CMIP6 will exercise the model in a variety of aspects. This, in turn, will provide the broad community with the ability to identify the strengths and weaknesses of CESM2. The participation to CMIP6 can therefore be seen as an extraordinary opportunity to document this new version, akin to the CESM1-CAM5 participation in CMIP5. All simulation results will be quickly post-processed to the standard CMIP format (owing to the improved CESM workflow designed by J. Dennis' group in the NCAR Computational Information Systems Laboratory in collaboration with the CESM Software Engineering Working Group, see Data Management section). The CMIP-formatted model output will therefore be quickly available for analysis by the University Community using the recently available NSF-funded Computational Information Systems Laboratory (CISL) CMIP Analysis Platform (https://www2.cisl.ucar.edu/resources/cmip-analysis-platform) and also for download from the Earth System Grid Federation (i.e. available to anyone).

In addition to the CESM2 standard 1° model, we will use CESM2-WACCM6 with enhanced middle and upper atmosphere representation (at 1° atmosphere and ocean horizontal resolutions) to simulate a variety of fields (ozone, volcanic aerosols, nitrogen deposition) needed to drive the CESM2-CAM6 version of the model. Because CESM2-WACCM6 is a different version then CESM2-CAM6 (higher top, explicit simulation of the quasi-biannual oscillation and interactive chemistry, including volcanic aerosols), it is necessary for CESM2-WACCM6 to perform the

DECK and historical simulations and computational resources are requested for those.

Overall, the CMIP6 simulations with the 1° version of the model will amount to approx. 239M Yellowstone core-hours. An additional 10M Yellowstone core-hours are planned for the spinup of ocean biogeochemistry and land carbon pools (these are additional steps to the overall procedure of building a pre-industrial control) and specific paleoclimate dataset creation (see PaleoWG request in Supplemental Material).

MIP acronym	MIP name	Name of primary sponsor(s)
AerChemMIP	Aerosols and Chemistry Model Intercomparison Project	Lamarque/Emmons
C4MIP	Coupled Climate Carbon Cycle Model Intercomparison Project	Lindsay
CFMIP	Cloud Feedback Model Intercomparison Project	Medeiros/Kay (CU)/Klein (LLNL)
DAMIP	Detection and Attribution Model Intercomparison Project	Tebaldi/Arblaster (Monash U.)
DCPP	Decadal Climate Prediction Project	Danabasoglu/Meehl
GeoMIP	Geoengineering Model Intercomparison Project	Tilmes/Mills
GMMIP	Global Monsoons Model Intercomparison Project	Fasullo/Kinter (COLA)
HighResMIP*	High Resolution Model Intercomparison Project	Neale/Bacmeister
ISMIP6	Ice Sheet Model Intercomparison Project for CMIP6	Lipscomb (LANL)/Otto-Bliesner
LS3MIP	Land Surface, Snow and Soil Moisture	D. Lawrence
LUMIP	Land-Use Model Intercomparison Project	D. Lawrence/P. Lawrence
OMIP/OCMIP	Ocean Model Intercomparison Project	Danabasoglu
PMIP	Palaeoclimate Modelling Intercomparison Project	Otto-Bliesner/Brady
RFMIP	Radiative Forcing Model Intercomparison Project	Gettelman/Neale
ScenarioMIP	Scenario Model Intercomparison Project	Meehl/O'Neill/P. Lawrence
VolMIP	Volcanic Forcings Model Intercomparison Project	Mills/Otto-Bliesner
Data only		
CORDEX*	Coordinated Regional Climate Downscaling Experiment	Mearns/Gutowski
DynVar	Dynamics and Variability of the Stratosphereâ€ Troposphere System	Marsh
SIMIP	Sea-Ice Model Intercomparison Project	Bailey/Holland/Jahn (CU)/Hunke (LANL)
VIAAB	VIA Advisory Board for CMIP6	Mearns/O'Neill

Table 1. List of CMIP6 Model Intercomparison Projects (MIPs) that CESM for which plans to perform the Tier 1 simulations. The asterisks (for HighResMIP and CORDEX) cells indicate the MIPs for which only partial participation is planned. The sponsors (right column) indicates the scientists (NCAR or otherwise listed) who will be responsible for ensuring that 1) simulations are correctly performed, 2) data are correctly posted and 3) analysis is performed. Data only MIPs are only requesting specific output streams but no additional simulations.

CMIP6 and the CESM2 release represent the core commitment of this CSL request. That being said, the forefront of computational climate science lies at high horizontal resolution. The aspect of high-resolution is aligned with the NCAR Strategic Plan (Imperative 3) since it bridges the weather-climate interfaces by providing many more explicitly resolved scales. Consequently, we are proposing to carry out some exploration in this area in the coupled ESM context and requesting a relatively small allocation (approx. 46M Cheyenne core-hours) for performing the pre-industrial control simulations for the high-resolution version of CESM2 (1/4° atmosphere); in addition, AMIP (prescribed sea-surface temperature) for 2 MIPs (GMMIP and HighResMIP) are requested. The size and scope of this request reflects 1) the very large cost of this version and 2) the limited experience of this version in a coupled ESM framework that would warrant extensive simulations. Additional time for development and testing of this version is also included within working group requests.

In terms of disk space and storage, using the current estimate for the CMIP6 DECK + Tier 1 (see [http://clipc-services.ceda.ac.uk/dreq/tab01_1_1.html)](http://clipc-services.ceda.ac.uk/dreq/tab01_1_1.html) and our CMIP5 experience (for which there was approximately a factor 6 between the raw data and the CMIP-processed posted data), we expect the CMIP6 1° simulations to generate approximately 2 PB. This is quite similar to the 50 GB/year generated during the Large Ensemble simulations. In addition, the high-resolution CESM2 simulations for CMIP6 are expected to generate will lead to approx. 270 TB.

Community Projects

Over the last 2 cycles of CSL proposals, we have defined a collection of experiments as Community Projects since they represent large simulations that are of interest to multiple Working Groups. Examples include the Large Ensemble or CESM2 development coupled simulations. Following the same philosophy, we have defined 5 projects for the current request, ranging from 3.5 to 46 M core-hours. The process for selection included a call for proposal, followed by a review of feasibility and cross-Working Group interest. The selected topics are

1. High-resolution CMIP6 simulations (42.5M core-hours): this request covers all the CMIP6 simulations to be performed under this computer allocation using the 1/4° version of CESM2. It is limited to the pre-industrial control, in addition to two AMIP (specified sea-surface temperatures) simulations. It will provide information on the behavior (e.g., modes of variability, biases, …) of this model that is currently unavailable.
2. WACCM simulations of future (2015-2100) atmosphere (3.5M core-hours): these simulations are aimed at adding 2 ensemble members to the current single simulation being performed under the CMIP6 allocation. Since WACCM will be providing the chemical fields for the other CESM simulations, these additional ensemble members will provide some information on the importance of internal variability on those fields.
3. Holocene (14.2M core-hours): The *transient Holocene* simulation will provide model data to more fully explore multidecadal and longer variability of, for example: ENSO and other modes of climate variability; monsoons and droughts; the AMOC; and tropical/extratropical linkages. This is an unprecedented CESM transient simulation covering the period from 9000 years ago until present.
4. PaleoSTRAT (5.4M core-hours): Two 100-year WACCM simulations will be carried out to evaluate the effect of aerosols, including volcanic injections, in the "last millennium" (focusing on 1750-1850 to include large volcanic eruptions such as Laki or Tambora). This is part of a project in collaboration with D.

Barriopedro, N. Calvo and R. Garcia (U. of Madrid, Spain), G. Chiodo (Columbia U.), and R. Neely (U. of Leeds, UK).

5. Ocean hindcast (17.4M core-hours): this request is for performing a 58-year forced ocean – sea-ice hindcast simulation for the 1958-2015 historical period with the 0.1° horizontal resolution version of the model, using the new JRA-55 forcing data sets. This simulation will provide unmatched statistics on ocean circulation.

Working Group Research Objectives and Requests

In this section, we describe in more details the overall research objectives specific to each working group. In addition, we provide for each WG the requested computing allocation, split between development and production and Years 1 and 2 (identified as D1, D2, P1 and P2, respectively). All numbers in the table below are in thousands of Cheyenne core-hours. More detailed information is available in the Supplementary Material document, which provides the full description of each WG request.

	Year 1		Year 2		Total per WG		
Working group	Prod	Dev	Prod	Dev	Year 1	Year 2	Y1+Y2
AMWG	5039	14507	10793	14839	19546	25632	45178
BGCWG	6765	8855	8215	9272	15620	17487	33107
CHWG	3626	2874	4510	3018	6500	7528	14028
CVCWG	20449	0	23096	0	20449	23096	43545
LIWG	4130	3521	6472	3738	7651	10210	17861
LMWG	7153	5976	7834	6643	13129	14477	27606
OMWG	6050	6862	6165	8020	12912	14185	27097
PaleoWG	10346	7817	12452	9435	18163	21887	40050
PCWG	3469	3629	4559	4438	7098	8997	16095
SDWG	6515	2066	4726	4034	8581	8760	17341
SEWG	0	8000	0	8000	8000	8000	16000
WAWG	8617	6352	12813	11483	14969	24296	39265
Total					152618	184555	337173

Community Projects							
High-resolution ocean hindcast	6000		11400				17400
WACCM 2-member ensemble	0		3485				3485
PaleoSTRAT	0		5395				5395
CESM High res	32000		10500				42500
Holocene	0		14206				14206
Total	38000		44986				

Total (dev+prod+comm)					190618	229541	420159
Target					190000	230000	420000

For all requests, the choice was left to the WG to balance between simulation throughput and cost (see Model Performance section below). Consequently, different estimates can be found for the same model configuration. In addition, the estimates for new versions have relied on simple scaling arguments (with number of level, tracers or horizontal resolution relative increases). Finally, the cost of running the same configuration on Cheyenne was estimated as being 0.82 the cost on Yellowstone, following the CISL recommendation.

Atmosphere Model Working Group (AMWG)
D1: 14.5M; D2: 14.8M; P1: 5.0M; P2: 10.8M; Total: 45.2M

The atmosphere model working group utilizes CSL resources primarily for the development of the Community Atmosphere Model (CAM) and associated capabilities. This encompasses the advancement of both the representation of the unresolved physical processes in parameterization schemes and the dynamical core processes, including tracer transport. It is also covers sensitivity experiments aimed at understanding the many interactions among the represented physical and dynamical processes across climate regimes and multiple timescales. The ongoing major changes in physics from CAM5 to CAM6 has opened up a variety of research and development areas for the AMWG that require significant computational resources to exploit, while making a concerted push towards higher resolutions that will define the frontier of climate science. The AMWG will also advance high resolution (both horizontal and vertical) and regional modeling capabilities through this CSL cycle. It will take the form of both improved global uniform high-resolution simulations, and regional refinement capabilities with a global configuration available through both the Spectral Element (SE) and Model Prediction Across Scales (MPAS) dynamical cores. This will inevitably require research with the existing CAM6 and future physical parameterizations in order to make them scale aware. Properties of scale-awareness enable schemes to work consistently between global high- and low- resolution grids and within regionally refined simulations where grid scales can vary by up to an order of magnitude. Finally, there will be an enhanced investment of resources in model assessment and validation through more non-standard techniques, akin to weather forecasting.

Biogeochemistry Working Group (BGCWG)
D1: 8.9M; D2: 9.3M; P1: 6.8M; P2: 8.2M; Total: 33.1M

The goal of the biogeochemistry working group is to produce a state-of-the-art Earth System Model for the research community that includes terrestrial and marine ecosystem biogeochemistry. While the magnitude of contemporary ocean uptake of anthropogenic carbon is constrained by observations to within 10%, the future uptake is uncertain. A primary objective of the BGCWG request is to estimate this future ocean uptake using CESM. Current research suggests that terrestrial ecosystems are at present a net carbon sink, but this conclusion masks considerable complexity and uncertainty with respect to future behavior. The availability of nitrogen, as well as other nutrients (e.g., phosphorus), alters the magnitude of the carbon cycle-climate feedback. Additional processes associated with ozone deposition and methane emission will alter the magnitude of the biogeochemical-climate feedbacks. Human activities from land use and land cover change play a very direct role in terrestrial ecosystem dynamics. The ambiguities in the mechanisms controlling the land carbon sink and their climate sensitivities translate into large uncertainties in future atmospheric CO_2 trajectories and climate change rates. Another primary objective of the BGCWG is to analyze these, and other, terrestrial feedbacks using CESM.

Chemistry-Climate Working Group (CHWG)
D1: 2.9M; D2: 3.0M; P1: 3.6M; P2: 4.5M; Total: 14.0M

The goal of the Chemistry-Climate Working Group is to continue development of the representation of chemistry and aerosols in CESM and to further our understanding of the interactions between gas-phase chemistry, aerosols and climate. The scientific motivation for these developments is the need to understand present-day and future air quality, to understand the role of climate change on tropospheric composition and changes in ozone in the lower stratosphere. The representation of tropospheric chemistry and aerosols continues to be developed and improved in CESM by the CHWG. Inorganic nitrate aerosols are being added within the framework of the Modal Aerosol Model (MAM4) using the MOSAIC (Model for Simulating Aerosol Interactions and Chemistry) treatment of aerosol thermodynamics, phase state and dynamic gas-particle mass transfer and heterogeneous chemistry. The formation and removal of secondary organic aerosols (SOA) will continue to be developed and evaluated as CESM evolves and more observational data sets from recent field campaigns become available. CAM-chem is a valuable tool for the interpretation of observations, and simulations with the improved nitrate and SOA schemes will be used to analyze recent campaigns. The previously developed very short-lived (VSL) organic halogen chemical mechanism will be used in model evaluations with field campaigns over remote oceans. The coupling of biogenic and fire emissions of chemical compounds and aerosols generated in the land model to the chemistry in the atmosphere will be evaluated and further developed in CESM2. The spectral element and MPAS dynamical models will provide valuable opportunities to study atmospheric chemistry, air quality and climate interactions on regional and local scales, and provide interpretation of field campaigns. As soon as large numbers of tracers can be transported efficiently in these models, the detailed tropospheric chemistry schemes will be tested in them. The CHWG will work with the Whole Atmosphere Working Group to perform the community simulations for DECK and CMIP6 and plan to provide simulations from a single model combining the full altitude range of WACCM with the full tropospheric and stratospheric chemistry scheme of CAM-chem. CAM-chem simulations will continue to be provided for other international model intercomparison and assessment activities, such as the World Meteorological Organization (WMO) 2018 ozone assessment.

Climate Variability and Change Working Group (CVCWG)
D1: 0; D2: 0; P1: 20.4M; P2: 23.1M; Total: 43.5M

The goals of the Climate Variability and Change Working Group are to understand and quantify contributions of natural and anthropogenically-forced patterns of climate variability and change in the 20th and 21st centuries and beyond by means of simulations with the CESM and its component models. With these model simulations, researchers will be able to investigate mechanisms of climate variability and change, as well as to detect and attribute past climate changes, and to project and predict future changes. The CVCWG simulations are motivated by broad community interest and are widely used by the national and international research communities. The highest priority for the CVCWG simulations is given to simulations that directly benefit the CESM community. The main focus over the next two years will be simulations intended for submission to CMIP6 including numerous "MIPs", lengthy control integrations with hierarchical configurations of CESM2, and AMIP and "Pacemaker" style historical runs. The CVCWG will contribute to the Detection and Attribution Model Intercomparison (DAMIP), Scenario MIP (ScenarioMIP), Flux Anomaly Forcing MIP (FAFMIP), and Cloud Forcing MIP (CFMIP). Analyses will target forced climate changes and associated uncertainties due to natural variability (assessed by running large ensembles), changes in variability and extremes and associated uncertainties, and changes across collections of ensemble members with different scenarios to assess forcing-related uncertainties.

Land-Ice Working Group (LIWG)
D1: 3.5M; D2: 3.7M; P1: 4.1M; P2: 6.5M; Total: 17.9M

The land-ice model capability has now reached a level of maturity that warrants significant computational investment. The first application of the requested resources will be to continue development of a self-consistent pre-industrial coupled ice-sheet/climate state, which will form the basis for future transient simulations. This spin-up will be computationally expensive and also employ novel component set combinations, in light of the 10^4-year Greenland Ice Sheet (GrIS) equilibration time scale due to characteristic mass balance and ice velocity. Additional resources will be required to perform mid-spin-up re-calibrations, based on validations to available observations. Once spin-up and validation/calibration exercises are complete, the group will perform a series of transient past and future coupled ice-sheet/climate simulations, such as the deglaciations during the Pliocene, the Last Interglacial (LIG) and the Holocene. Future simulations will address the fully coupled response of the Greenland ice sheet to anthropogenic forcing on CMIP6 and longer time-scales.

Land Model Working Group (LMWG)
D1: 6.0M; D2: 6.6M; P1: 7.2M; P2: 7.8M; Total: 27.6M

The goals of the Land Model Working Group are to continue to advance the state of the art in modeling Earth's land surface, its ecosystems, watersheds, and socioeconomic drivers of global environmental change, and to provide a comprehensive understanding of the interactions among physical, chemical, biological, and socioeconomic processes by which people and ecosystems affect, adapt to, and mitigate global environmental change. Land biogeophysical and biogeochemical processes are intimately linked and therefore it is not possible to separate land biogeophysics development from land biogeochemistry development. For this and previous allocation requests, land biogeochemistry model development has been included in the Land Model Working Group request. A portion of the proposed terrestrial carbon cycle production work has been included in the Biogeochemistry Working Group request. The Land Model Working Group has pursued an ambitious program of model development, which will culminate with the release of CLM5 during this CSL allocation period. Several additional large development projects have been progressing in parallel to CLM5 development including a multi-layer canopy scheme, a hill-slope hydrology model, and the Ecosystem Demography version of CLM. These projects will continue into the next CSL along with other development projects. Parameter estimation/calibration is an increasingly important feature of CLM development. In addition, land processes and their role in climate variability and change have gained significant expanded focus in CMIP6. Land-focused MIPs within CMIP6 include LUMIP (Land-use MIP), LS3MIP (Land surface, soil moisture and snow MIP), and C4MIP (Coupled Climate Carbon Cycle MIP). Together, these MIPs address the main feedbacks and forcings from the land surface, and also include a benchmarking land-only MIP ("LMIP", which is part of LS3MIP (see CMIP6 section).

Ocean Model Working Group (OMWG)
D1: 6.9M; D2: 8.0M; P1: 6.0M; P2: 6.2M, Total: 27.1M

The primary goals of the Ocean Model Working Group are to advance the state-of-the-science in the capability and fidelity of the CESM ocean component in support of specific science objectives of the broad CESM effort and community and to conduct curiosity driven research with CESM to advance our understanding of the role of the oceans in the Earth's climate system. Our overall objectives continue to be the leaders in new model developments, particularly in parameterizations, and to deliver a state-of-the-science ocean model to the CESM community for the next generation of the CESM model. The primary development activity of the OMWG for the next 2-3 years will be the incorporation of a new ocean model (dynamical core) within the CESM framework. As we look beyond CESM2, it is necessary to formulate a plan for the next generation ocean model component. The Parallel Ocean Program (POP) model has been used as the ocean component of CESM for more than a decade, but will not be developed further by the Los Alamos National Laboratory. There is an on-going process for a selection of the POP replacement, with a decision expected in October 2016. The second development goal is to complete ongoing parameterization development efforts that were started over the last few years. Additional resources are also requested for i) several data assimilation developments; ii) final evaluations of a new atmospheric data set used for forcing ocean – and sea-ice coupled simulations; iii) testing of this new data set for use in high-resolution version of the POP ocean model; and iv) developing and testing a regional ocean model with a biogeochemical model for the coral triangle region. Production request targets several science goals. These include coupled experiments with the recently-developed one-dimensional ocean model that would be used to decipher the role of ocean dynamics in climate and its variability;

evaluations of CESM-DART reanalysis in comparison with initialized, coupled hindcast simulations; and investigations of the Labrador Sea hydrographic properties, the North Atlantic Oscillation-related surface heat fluxes, and surface freshwater flux anomalies on the mean and variability in the North Atlantic, with a particular focus on the Atlantic meridional overturning circulation. We note that the OMWG will lead a proposed community project involving preparation and integration of a high-resolution ocean – sea-ice (both at 0.1° horizontal resolution) forced hindcast simulation. Use of such high-resolution (eddy-permitting / -resolving) ocean models in routine climate applications – requiring many long simulations – remains, however, prohibitively expensive and are not included in this proposal.

Paleoclimate Working Group (PaleoWG)
D1: 7.8M; D2: 9.4M; P1: 10.3M; P2: 12.5M; Total: 40.0M

The main development goal for the Paleoclimate Working Group is to provide the community with expanded capabilities in CESM for application to a wide range of paleoclimate research problems on multiple time scales and time periods. The PaleoWG plays a unique role in the CESM Community as it acts as a testbed for a wide range of forcings and ice-sheet physics. Examples include testing new configurations of CESM, such as the capability to simulate the inception and retreat of Greenland, North American, and Eurasian ice sheets when coupled to CESM and to test emission scenarios for a large asteroid impact with CARMA coupled to WACCM. Efforts are also being focused on development of a version of CESM2 for deep-time paleoclimate research. The production goal is to provide benchmark simulations of past climates to the community. These simulations offer the opportunity to test the CESM for various forcing conditions, carry out detection and attribution studies, and improve confidence in its application for the future. The working group carries out experiments as part of international intercomparison projects – CMIP6, PMIP4, VOLMIP, and ISMIP6. Our proposed Production simulations are the Tier 2 and 3 simulations of PMIP4 and VOLMIP, which have been proposed by these MIPs as a coordinated set of sensitivity experiments to complement and enhance understanding of the CMIP6 Tier 1 simulations.

Polar Climate Working Group (PCWG)
D1: 3.6M; D2: 4.4M; P1: 3.5M; P2: 4.6M; Total: 16.1M

The overall development objective for the Polar Climate Working Group is to ensure that CESM has state-of-the-art abilities to simulate polar climate in a time of rapid polar change. The CSL resources are used to facilitate the use of cutting edge observations and techniques (e.g., data assimilation, satellite simulators, high-resolution) by PCWG members towards our overall development goal. Here, resources are sought to incorporate new polar-relevant physics and diagnostics into the sea ice model (CICE) and atmospheric model (CAM) used in CESM, and to test a new ice-only high-resolution model configuration. The over-arching PCWG production goal is to enable important and topical polar science research using

CESM. Our proposed experiments leverage and enhance community production experiments and expertise. The production experiments will be run by a collection of university and NCAR investigators.

Societal Dimensions Working Group (SDWG)
D1: 2.1M; D2: 4.0M; P1: 6.5M; P2: 4.7M; Total: 17.3M

The Societal Dimensions Working Group seeks to improve the understanding of the interactions between human and earth systems by enhancing CESM and its application through studies of climate change impacts, adaptation, and mitigation that use CESM output in their analyses. The request supports core projects in linkages between CESM and integrated assessment models (IAMs), while also reaching out to engage additional impacts, adaptation and vulnerability user communities consistent with the broadened focus of the working group. It also supports CESM/SDWG contributions to important community processes such as the design of CMIP6 experiments, particularly those related to future scenarios and land use (ScenarioMIP, LUMIP). More specifically, development objectives will include further evolution of methods for coupling CESM with human system models and improving the implementation of land use, a key process through which human and earth systems interact. This request also supports production simulations that will apply CESM to analyses of regional climate and impacts, agricultural impacts, and geoengineering, while also supporting CESM contributions to important community projects such as CMIP6.

Software Engineering Working Group (SEWG)
D1: 8M; D2: 8M; P1: 0; P2: 0; Total: 16M

The present request enables the Software Engineering Working Group to coordinate the computational development of the CESM model components, oversee the evolving design of the CESM as new model components, new model grids and new model physics are added to the system and at the same time engineer the model system to obtain optimal throughput and efficiency. This continues to be particularly challenging as the number of model configurations, model complexity and model resolutions are rapidly increasing. Numerous tests are carried out for each CESM revision on all targeted production platforms to ensure required functionality (such as exact restart capability), correct results (such as bit-for-bit reproducibility where it is expected), tracking of memory and performance metrics (to determine if these have changed relative to the previous revision) and other key production requirements (such as optimizing performance of new revisions, especially where new component science has been introduced). In addition, this testing also ensures the robustness of the continuing and significant model infrastructure development, such as the improvements to changes to the model driver, coupler, tools, and scripts. The request is also needed to support the upcoming CMIP integrations as well as new workflow capabilities that will be associated with those integrations.

Whole Atmosphere Working Group (WAWG)
D1: 6.4M; D2: 11.5M; P1: 8.6M; P2: 12.8M; Total: 39.3M

The Whole Atmosphere Working Group research plan involves development designed to continue the move towards a unified sun-to-earth modeling (WACCM, Whole Atmosphere Community Climate Model) framework with high fidelity, and production runs for science and community projects. This involves continuing work on a number of development projects across NCAR laboratories and outside collaborators. The development request focuses on building a unified sun-to-earth modeling framework. This will include advancing the photolysis treatment, exploring higher vertical resolution, improving representation of gravity waves, and bring WACCM-X, the solar weather model, up to the same climate model version as the rest of CESM. This will provide a framework for the simulation of space weather within CESM, thereby taking advantage of the explicit representation of the full processes in the lower atmosphere that are affecting the upper atmosphere. On the production side, simulations using WACCM-CARMA will be used to look at aerosols in the upper troposphere/lower stratosphere region, and explore possible geo-engineering solutions. This provides reference simulations that can be used in the development of simplified aerosol models. The production allocation includes simulations that will provide contributions to the WMO2018 ozone assessment. Production will also include WACCM contributions to the CMIP DECK experiments for a specified chemistry version (WACCM6-SC, which is considerably less expensive than the full chemistry version used in the CMIP6 simulations and is a key tool for exploration of dynamical interactions) and ISA-MIP for intercomparison of stratospheric aerosols.

Data Management and Archive Requirements

As part of this CSL proposal, each working group generated estimates of the data volume associated with each proposed development and production experiment (listed in the Supplementary material). As in the previous CSL proposal, we have defined a CESM Data Management and Data Distribution Plan (see http://www.cesm.ucar.edu/management/docs/data.mgt.plan.2011.pdf) that has Production and Development data stored and distributed via different strategies, with each tailored to suit the different user needs. Note that the present estimate takes into account the recent significant improvement in data compression from the use of the netCDF-4 standard. This was made possible through a major refactoring of the simulation workflow that incorporates a parallel Python based utility to convert uncompressed history time-slice output data files to compressed netCDF-4 formatted variable time-series output data files as part of the CESM run script. All numbers are in Terabytes (TB).

Based on this current estimate of CMIP6 (which might still change owing to the fact that the model output requests for the different MIPS are not finalized yet), we are expecting to generate approximately 8.3 PB of data during the upcoming CSL cycle.

	Year 1		Year 2		Total per WG		
Working group	Prod	Dev	Prod	Dev	Year 1	Year 2	Y1+Y2
AMWG	32	39	85	99	71	184	255
BGCWG	76	250	92	240	326	332	658
CHWG	80	55	75	60	135	135	270
CVCWG	331	0	366	0	331	366	697
LIWG	42	49	77	41	91	118	209
LMWG	79	6	43	7	85	50	135
OMWG	166	276	138	114	442	252	694
PaleoWG	123	79	260	112	202	372	574
PCWG	42	49	63	60	91	123	214
SDWG	104	32	100	34	136	134	270
SEWG	0	0	0	0	0	0	0
WAWG	294	137	311	190	431	501	932
Total					2341	2567	4908
Community Projects							
High-resolution ocean hindcast	44		84				128
WACCM 2-member ensemble	0		88				88
PaleoSTRAT	0		200				200
CESM High res	200		70				270
Holocene	0		384				384
Total	244		826				1070
Total: dev+prod+comm					2585	3393	5978
CMIP6							2270

a. Data Archiving

The total data volume expected from to be generated from development, production and community projects simulations amount to approximately 8.3 PB, with 2.3 PB associated with the CMIP6 portion of the CSL request alone. This represents a similar amount in data generation relative to the previous CSL proposal. To manage this data volume we will continue our current archival strategy:

Development: Output data will be primarily stored on the requested glade partition (see below). If necessary, minimal output will be written on the HPSS. We expect that this will account for approximately 10-20% of the development output. In addition, for a period of at most 36 months after creation, at which point it will be removed, unless retention is requested from the relevant working group co-chairs. "One-off" development experiments will be removed more quickly at the working group co-chair's discretion.

Production: Output data will be stored on the HPSS for a period of four years. It will then be gradually cut back to 50% of its initial volume over a period of three additional years, based on usage and anticipated demand. This data level will be maintained for three more years. Afterward, each working group will determine what data are to be removed and at what rate, as the archived data is gradually reduced to an acceptable level, as determined by data archiving costs at the time.

Archive Management: All the experiments listed in the proposal will make full use of the existing CESM Experiment Database (see http://csegweb.cgd.ucar.edu/expdb). This database has been developed over the last year and contains details about the run configuration and establishes provenance. The database application runs an automated monthly email reminder script triggered off dates stored in the database fields; as such, it will be used to remind all affiliated users with the experiment, including scientific leads and software engineers, to prune their data from HPSS according to the CESM Data Management and Data Distribution Plan.

b. Data Distribution:

Development: In general, output data will be made available only to the working group members that are directly involved in the experiments. For working group members that do not have access to the NCAR HPSS, these data will be made available via the Earth System Grid (ESG).

Production: Output data will be made available according to the guidelines established by the CESM Data Management and Data Distribution Plan, which is formulated by the CESM SSC, NCAR and NSF. Initially, access is restricted to the working group members directly involved in the experiments. After a period of no more than 12 months after creation, these data will be made available to the community via the ESG. The CMIP6 model output will be made available to the CMIP Analysis Platform and ESG in accordance to the CMIP protocol. It is expected that, owing to the improvement in the postprocessing workflow, the

c. Data Analysis and Visualization Request

The simulations produced under development and production CSL resources will require considerable analysis and visualization. For these needs, we request access to the Geyser and Caldera data analysis and visualization (DAV) clusters. This will require standard interactive access to these clusters for the working group members that have CSL HPC resources and for additional participants who are helping in the analysis of these simulations. Currently this includes about 150 participating scientists but is subject to change with changing working group members and involvement.

d. GLADE project file space (total request: 2.5PB)

In order to minimize the usage of HPSS for storing development results, we request 1.5 PB of CESM GLADE project space. This will also enable efficient access to highly utilized CESM simulation output and forcing data used in coupled integrations. It will also allow for the post- processing of community project integrations. This is significantly larger than the current CESM project space, but we believe is necessary given the increased number of integrations and simulation output that will be performed under this proposal. This space is collectively managed by the CESM working groups. The current allocation (table below) is clearly indicating a highly used and useful capability.

Space	Used	Quota	% Full
/glade/p/cesm	208	228	91
/glade/p/cesm/amwg_dev	42	42	99
/glade/p/cesm/bgcwg_dev	65	66	99
/glade/p/cesm/chwg_dev	76	84	91
/glade/p/cesm/liwg_dev	3	5	69
/glade/p/cesm/lmwg_dev	16	17	95
/glade/p/cesm/omwg_dev	90	91	98
/glade/p/cesm/palwg_dev	114	120	95
/glade/p/cesm/pcwg_dev	40	40	99
/glade/p/cesm/sdwg_dev	50	57	88
/glade/p/cesm/wawg_dev	91	91	100
/glade/p/cesmLE	124	125	99
/glade/p/cesmLME	79	125	64
/glade/p/cesm0005	179	253	71
Total	590	731	81

The increase in GLADE space is commensurate with the increase in computing power from the previous CSL proposal.

In addition, because of the need to postprocess the raw CESM output into the CMIP format (Figure 1), we make an additional request of 1 PB to have sufficient glade space to allow for multiple simulations to be performed in parallel and an efficient conversion. Indeed, the CESM community projects "Large Ensemble" and "Last Millennium Ensemble" have greatly benefited from the allocation of a significant specific allocation (500 TB) of GLADE project space. Original model output from these two sets of runs amounted to about 2 PB, but the existence of ample GLADE

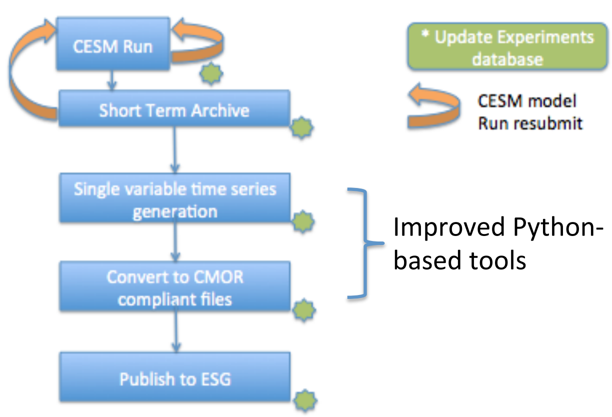

Figure 1. Schematic of the workflow designed to facilitate and greatly reduce the required time between the CESM simulation and the publishing to ESG. The use of Python-based (and parallelized) has reduced the time-series generation step by one or two orders of magnitude.

project space allowed CESM to analyze and postprocess CESM output, for multiple completed simulations, without the use of the HPSS as "overflow" space during a model run, and without the concern of output being scrubbed. This GLADE space allowed for much more efficient use of HPSS, without output being written needlessly to HPSS, retrieved at a later time for post processing, and then the unneeded data removed from tape. A few hundred TB of glade project space significantly improved the workflow of these two large community projects, and reduced the impact on HPSS and consumption of tapes. This proposal therefore continues to use and expand this improved workflow process.

Model Performance

The main configuration targeted in the CESM2 release will have the 1° resolution for all components (except land-ice, which uses its own internal mesh refinement). It is critical to have a good understanding of the overall performance (cost and throughput) for this configuration. Using the most recent simulations performed with CESM2 on Yellowstone, we see that this version of the model is giving a very noteworthy throughput of approximately 18 simulated years per wall-clock day. Note that this version has the finite volume (FV) dynamical core, 32 levels, and only the standard Modal Aerosol Model (MAM4; Liu et al., 2016) chemistry.

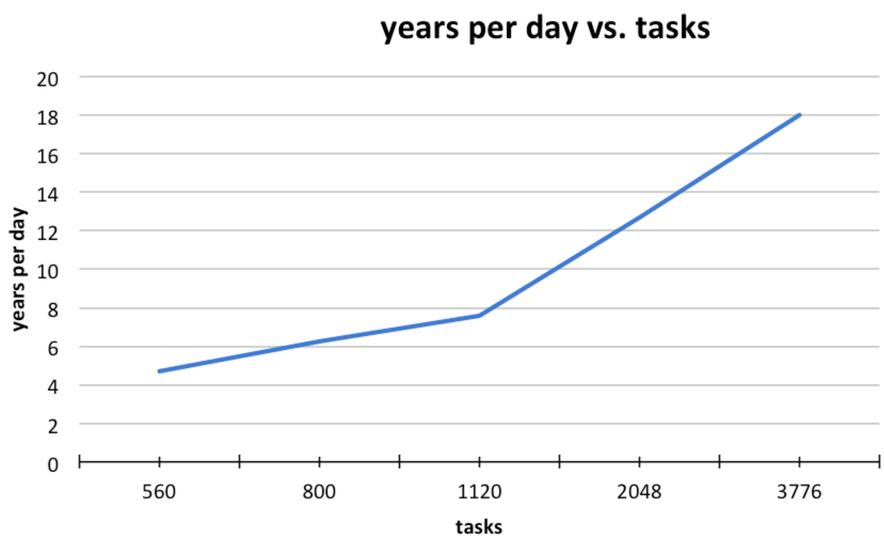

Such throughput is necessary to enable the performance of numerous simulations (as proposed here) in a reasonable amount of time (such as the 1-year window for the CMIP6 simulations on Yellowstone). This, obviously, comes with an increased cost (approx. 5K/year for 3776 tasks vs 1K/year for 560 tasks). A significant portion of this performance was obtained by optimizing the layout of each component. However, a significant amount of effort has also been applied to improving the level of vectorization within CESM. Based on detailed performance analysis of the most expensive 10-20 modules within CESM, it was decided to initially target the CAM model for optimization. A number of computational modules within the physics of CAM have been optimized and reintegrated back into the CESM2 release code base; much of this work has been done with the support of the

CISL Application Scalablity and Performance (ASAP) group. These optimizations include restructuring of code call structure to improve vectorization, elimination of assumed shape arrays, elimination of elemental functions, and improvements to variable alignment. One of these modules (the Morrison-Gettelman microphysics package version 2, MG2), was optimized and its cost reduced by approximately 50% on multiple Intel platforms. A description of this effort has been documented as a NERSC case study for code optimization (He, et al. 2016). As of June 2016, approximately 6 modules have been optimized and reintegrated back into the CESM code base and have reduced the cost of CAM by approximately 14%. Several additional modules are still being optimized but have not been fully reintegrated back into the code base. We expect an additional reduction of 10% reduction in the version of CAM released in CESM2.

Clearly, the possibility of performing any number of high-resolution simulations will necessitate a highly optimized code to reduce as much as possible its very high cost and to increase throughput. For that purpose, an effort lead by Dr John Dennis of the ASAP group has made significant progress optimizing the spectral element (SE) dynamical core used with in CAM for the ¼° (ne120) resolution. This effort has been funded as an Intel Parallel Computing Center for Weather and Climate Simulations (IPCC-WACS) and has been further supported under the NERSC Exascale Science Application Program (NESAP). IPCC-WACS and NESAP projects have been focused on increasing thread, task and instruction parallelism inherent in the code. Depending on the science objectives, the SE dynamical core can consume 30-88% of the total cost of CAM. A critical limitation in the previous version of CAM-SE was that it was not possible to scale beyond a single element per hardware-core. This single element per hardware-core limitation has recently been eliminated by a redesign of the OpenMP threading implementation within the dynamical core. These threading changes along with improvements to both the code vectorization and L3 cache behavior has significantly reduced the cost to execute the SE dynamical core at large core counts.

The impact that these changes have on the scalability of the SE dynamical core is shown in Figure 2 using a reduced problem size running on Yellowstone. Because the SE dynamical core has excellent weak scaling characteristics, we expect these 900 to 21,600 core results to be a good predictor of NE120 performance. For the purpose of this test, the SE dynamical core, which is also known as High Order Methods Modeling Environment (HOMME), is configured at a resolution of ne=30, and choose core counts such that the number of spectral elements per hardware core ranges from 6 element/HW-core (900 cores) to ¼ elements/HW-core (21,600 cores). Note that the target resolution of ne=120 will have 16-times the amount of parallelism as the ne=30 testcase. Also note that HOMME in this case has 26 vertical levels and 25 tracers, which is similar to a standard CAM6 configuration.

Relative cost (which indicates the cost of the dynamical core relative to the previous version of HOMME using 6 spectral elements per hardware core, 900 HW-cores) is also included as dotted blue and red lines. We use 6 spectral elements per hardware core as a base cost because CAM is typically run using approximately 4-6 elements per hardware core. For example, the ASD simulations (Small et al., 2014), which were performed on Yellowstone in 2012-2013 used a ¼ degree CAM and utilized 21,600 cores, had a per-core problem size of 4-elements per core. These cost lines allow for an easy comparison of the impact that these optimizations may have on

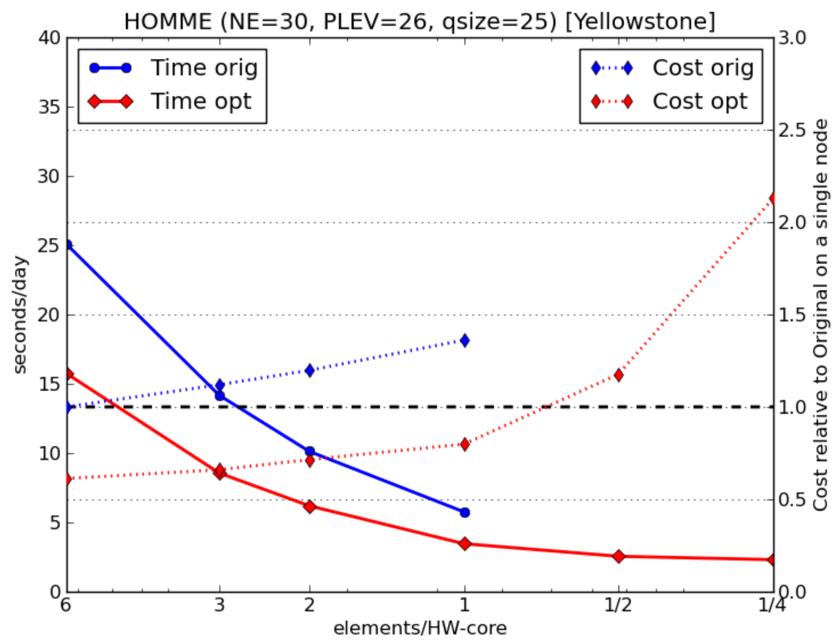

Figure 2. The solid blue line (labeled orig) corresponds to a version of the SE dynamical core used in previous versions of CESM before the IPCC and NESAP funded optimization effort began. The solid red line (labeled opt) optimized versions of the SE dynamical core. A significant reduction in execution time for the dynamical core is apparent for the opt versions of HOMME.

simulation cost. Specifically we take advantage of these new optimizations by either reducing the cost of the simulation at a fixed core count, or increase the number of cores that are used for a similar cost. For example, if we draw a vertical line from the cost to run HOMME using 6 elements per hardware core it is apparent that the cost of the 'opt' version of the code reduces the cost of the SE dynamical core by approximately 40%. If we draw a horizontal line from the cost to run HOMME using 6 elements per hardware core to the dotted red line, it is apparent that we can utilize ~12 times the amount of hardware cores with only a 19% increase in the computational cost using a SE dynamical core based on the opt code base.

Additional optimization of other CESM2 configurations (such as WACCM) is ongoing, to ensure maximal use of the requested computer allocation and sufficient throughput (especially for the CMIP6 simulations).

Summary

Earth System models are the most powerful tools for meeting the intellectual challenge of understanding the climate and the Earth system: they are the only scientific tool capable of integrating the myriad physical, chemical and biological processes that determine past, present and future climate. They are also essential for synthesizing, through data assimilation, diverse in-situ and remotely sensed observations. They are critical tools for testing hypotheses, confirming understanding and for making predictions of use to society and policy makers. The development of CESM is unique in that it occurs through strong partnership with scientists from universities, national laboratories, and other research organizations. CESM enables the investigation of new scientific problems through partnerships with a community broader than ever before, and it is now enabling many new partnerships including those involved in adaptation and mitigation research.

While previous model versions have been critical tools for advancing our understanding of earth system processes, challenges remain and new science question arise. The development of CESM2 is almost finalized and has been a focus of all working groups for the last 2 or more years. This development includes advances in new physical, chemical, biological and biogeochemical processes, along with high-resolution capabilities. High priority production simulations include those that allow the broad scientific community to participate in coordinated international modeling activities as well as benchmark simulations to document CESM components and coupled system functioning. More specifically, the participation of CESM2 to CMIP6 will provide a wealth of information on the model qualities and shortcomings. This will provide the CESM community with the necessary knowledge to make adequate use of CESM2, in addition to providing a partial roadmap towards CESM3. Additionally, emerging priorities for climate prediction at regional scales requires new production and development efforts to further understanding and simulation fidelity.

The CSL computer resource remains indispensable to carry out this ambitious agenda. The objectives and priorities outlined in this proposal emanate directly from the community of scientists who participate in the CESM project through the 12 working groups and the CESM Scientific Steering Committee. They were developed, refined and prioritized after a several month process with the goal of producing a coherent and coordinated plan for the use of the CSL resource over the upcoming period of performance. All of the proposed experiments will fill important development and production needs and contribute to the overall project priorities. However, note that in the individual requests present in the appendices, simulation experiments have been prioritized in the event that a reduced allocation is awarded. The work outlined requires a total of 250 million Yellowstone core-hours over the period November 2016 through December 2017; this allocation will mostly focus on performing CMIP6 simulations with the 1° version of CESM2. In addition, we are requesting a total of 420M Cheyenne core-hours to carry all other tasks, development and production, including Community Projects. We are projecting to generate more than 8 PB of data, of which less than 5 PB are expected to be written on the HPSS. This will be achieved through our continued (and highly valuable and successful) use of GLADE as the primary storage space for all development simulations. Additional GLADE space is requested for the intermediate storage of the CMIP6 data; this will be key to enable a fast postprocessing of the raw CESM data into the CMIP format, which is then posted on the ESG. In total, the GLADE request is for 2.5 PB. This level of computational and storage will allow 1) CESM2 to be finalized for its release in early 2017, 2) to perform CMIP6 simulations that are of interest to the CESM Community and provide diagnostics on model performance in all aspects of an Earth System, 3) explore cutting-edge high-resolution configurations and 4) start the development towards CESM3.

CESM gratefully acknowledges our primary sponsors,
The National Science Foundation and
The Department of Energy

References

Bacmeister, J. T., M. F. Wehner, R. B. Neale, A. Gettelman, C. Hannay, P. H. Lauritzen, J. M. Caron, and J. E. Truesdale (2014), Exploratory highresolution climate simulations using the Community Atmosphere Model (CAM), J. Clim., 27, 3073–3099.

Eyring, V., Bony, S., Meehl, G. A., Senior, C. A., Stevens, B., Stouffer, R. J., and Taylor, K. E.: Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) experimental design and organization, Geosci. Model Dev., 9, 1937-1958, doi:10.5194/gmd-9-1937-2016, 2016.

He, H. C. Kerr, J. M. Dennis CESM case study, http://www.nersc.gov/users/computational-systems/cori/application-porting-and-performance/application-case-studies/cesm-case-study/

Knutti, R., D. Masson, and A. Gettelman (2013), Climate model genealogy: Generation CMIP5 and how we got there, Geophys. Res. Lett., 40, 1194–1199, doi:10.1002/grl.50256.

Liu, X., Ma, P.-L., Wang, H., Tilmes, S., Singh, B., Easter, R. C., Ghan, S. J., and Rasch, P. J.: Description and evaluation of a new four-mode version of the Modal Aerosol Module (MAM4) within version 5.3 of the Community Atmosphere Model, Geosci. Model Dev., 9, 505-522, doi:10.5194/gmd-9-505-2016, 2016.

Small, R. J., et al. (2014), A new synoptic scale resolving global climate simulation using the Community Earth System Model, J. Adv. Model. Earth Syst., 6, 1065–1094, doi:10.1002/ 2014MS000363.

Wehner, M. F., K. A. Reed, F. Li, Prabhat, J. Bacmeister, C.-T. Chen, C. Paciorek, P. J. Gleckler, K. R. Sperber, W. D. Collins, A. Gettelman, and C. Jablonowski (2014), The effect of horizontal resolution on simulation quality in the Community Atmospheric Model, CAM5.1, J. Adv. Model. Earth Syst., 6, 980–997, doi:10.1002/2013MS000276.