CESM

Community Earth System Model

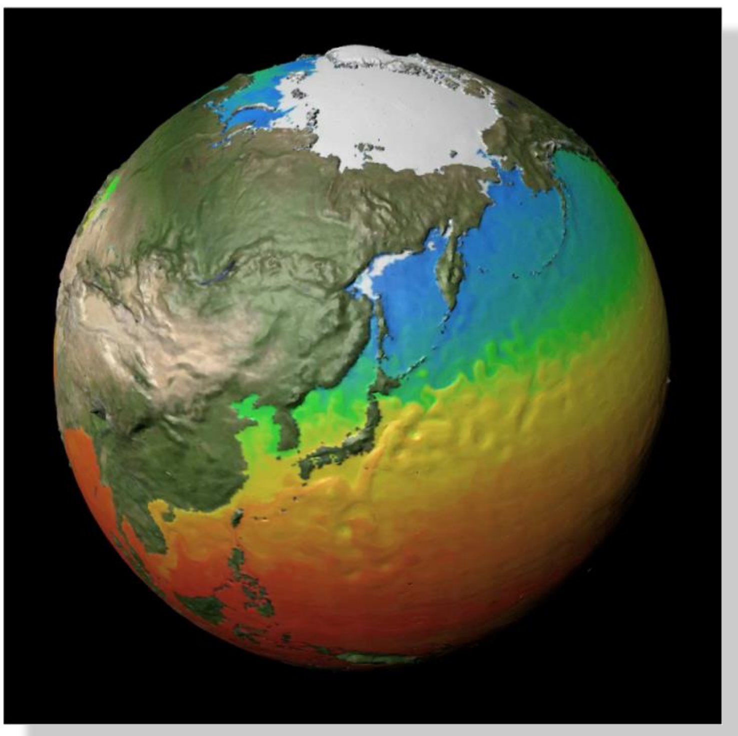

Proposal for CSL Resources

PI: Gokhan Danabasoglu, CESM Chief Scientist

Period of Performance: 01 November 2018 – 31 October 2020

Total Request: 460 M Cheyenne Core-Hours

Table of Contents

Cover image courtesy of J. Small and T. Scheitlin (NCAR)

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

Introduction

The Community Earth System Model (CESM) project is a community effort that requires collaborations 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 the Climate Simulation Laboratory (CSL) computational resources. The CESM and its predecessor, the Community Climate System Model (CCSM), have been at the forefront of both national and international efforts to understand and predict the behavior of Earth's climate. Output from numerous simulations using CCSM and CESM are routinely used in thousands of peer-reviewed studies to better understand the processes and mechanisms responsible for climate variability and change. Specifically, the overwhelming majority of these studies make use of CCSM's and CESM's contributions to the Coupled Model Intercomparison Project phase 3 and phase 5 (CMIP3 and CMIP5). Evaluation of the first generation CESM, CESM1, has ranked it among the best climate models in the world (Knutti et al. 2013). As a further testament to this point, the two primary manuscripts, Gent et al. (2011) and Hurrell et al. (2013), that introduced and described the previous two versions of the model, CCSM4 and CESM1, have been cited > 1550 and > 650 times, respectively, since their publications. Significant CSL-supported efforts such as the CESM Large Ensemble (CESM-LE; Kay et al. 2015), its sibling CESM Decadal Prediction Large Ensemble (CESM-DPLE; Yeager et al. 2018), and the CESM Last Millennium Ensemble (CESM-LME; Otto-Bliesner et al. 2016) have been key in advancing our understanding of the climate system and its variability and predictability, supplementing CESM's contributions to the CMIPs with community-driven science efforts. CESM source code and simulation output are made freely available to the broad scientific community. Additionally, CCSM and CESM simulations provide the quantitative modeling basis for 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 the National Science Foundation (NSF) and the Department of Energy (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, seasonal-to-decadal climate predictions, and projections of future climate change.

A major milestone that was accomplished during our current allocation cycle was the long-waited, community release of the second-generation CESM, CESM2, specifically the release of CESM2.0 in early June 2018. This new version contains many substantial science and infrastructure improvements and capabilities for use of the broader CESM and international community. These new advancements include: an atmospheric model component that incorporates significant improvements to its turbulence and convection representations, opening the way for an analysis of how these small-scale processes can impact the climate; improved ability to simulate modes of tropical variability that can span seasons and affect global weather patterns; a land ice sheet model component for Greenland that can simulate the complex way the ice sheet moves and does a better job of simulating calving of the ice into the ocean; a global crop model component that can simulate both how cropland affects regional climate, including the impacts of increased irrigation, and how the changing climate will affect crop productivity; a wave model component that simulates how wind creates waves on the ocean, an important mechanism for mixing of the upper ocean; an updated river model component that simulates surface flows across hill sides and into tributaries before entering the main river channel; and a new set of infrastructure utilities that provide many new capabilities for easier portability, case generation and user customization, testing functionality, and greatly increased robustness and flexibility. A full list of updates with more technical descriptions is available at http://www.cesm.ucar.edu/models/cesm2/whatsnew.html.

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 through the eleven[1] CESM working groups (WGs) and the CESM Scientific Steering Committee (SSC, whose membership consists of not only NCAR scientists but also scientists from universities and government laboratories). Specifically, to prepare this proposal, each WG has consulted with their constituents, beginning at the June 2018 Annual 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. During drafting of the WG plans, there were substantial discussions and planning across the WGs as well. The resulting drafts were further reviewed, revised, refined, and prioritized through a process of exchange across the different WGs, with the goal of producing a coherent and coordinated plan for the use of the CSL resources over the upcoming period of performance. The plans and resource requests of the individual WGs and community projects, which are included in the Supplementary Material, then served as the source material for further review by the CESM SSC. The responsibility of the SSC in this proposal was to review the overarching development and production simulations, including priorities, as well as the *community projects* that require significant resource sharing across WGs for the entire CESM project and to provide input on the main development and production activities and the required computing and data resources. A similar process was followed during the last several CESM CSL proposals and, we believe, has resulted in a coherent overview of the testing, development, and application needs of the broad CESM project.

[1] The recent staff reductions and other departures at NCAR had detrimental impacts on the CESM Societal Dimensions Working Group (SDWG) leadership, also with impacts on internal WG membership. The future of societal dimensions related work at NCAR and, in particular, within CESM is being discussed with no immediate clear paths. As such, this proposal does not include a request from the SDWG. However, many of the simulations that would be of interest to this WG are already included as part of either CESM DECK and Tier 1 simulations or in other WGs' requests.

Overarching Priorities

Fundamental scientific obstacles in the development of CESM2 that were not anticipated in the last CSL proposal required a nearly two-year extension in the development and simulation efforts. This extended activity was carried out with the full support of the SSC, CESM Advisory Board (CAB), funding agencies, and NCAR management and also was summarized at the midterm review of the last CSL proposal in September 2017. As detailed in the Accomplishments document, the extended period of development required all WGs to donate computational resources and substantial personal effort to addressing these shortcomings. A significant portion of the previous CSL proposal was predicated on the existence of a viable version of CESM2. Thus, as a result of the two-year extension in the development of CESM2, some of the major goals of this CSL request remain the same as in the last proposal. These goals are:

1. Perform CMIP6 (CMIP phase 6) DECK (Diagnostic, Evaluation, and Characterization of Klima) and MIPs (Model Intercomparison Project) Tier 1 simulations, primarily using the nominal 1° model version,
2. Perform many of the MIPs Tier 2 and Tier 3 simulations that are particularly designed to address specific science questions that arose from the previous MIP exercises,
3. Continue development and testing of the high-resolution configurations for CESM2,
4. Continue component model development efforts towards inclusion in the next generation model version, i.e., CESM3, as detailed in the WGs' requests, including development of the new ocean model, Modular Ocean Model version 6 (MOM6); re-evaluation of coupling methods between sea-ice and the ocean model; and incorporating new capabilities in the Community Ice Sheet Model (CISM) with the goal of supporting climate simulations with a dynamic Antarctic ice sheet.

As such, this CSL request has two primary thrusts. The first is a focus on applications and pushing the frontier of resolution with the recently released CESM2 (goals # 1, 2, and 3). The second is to drive the next generation model, whose development has continued even through the period of delay in CESM2 (goal # 4).

CESM's participation in CMIP6 DECK and MIPs simulations is its broadest community project. These simulations reach a vast national and international group of scientists and researchers who rely on NCAR, and CESM in particular, to perform these simulations. Among the modeling groups which contribute to CMIP6, CESM is very unique in its community involvement and the level of transparency with which it approaches model development. As indicated below, the requests for CESM's involvement in a particular MIP's simulations came directly from the broader user community and subsequently incorporated in this proposal. CESM's contributions to the USGCRP at a national and IPCC at an international level are realized through these CMIP simulations. Furthermore, many national and international researchers make use of CESM in their research proposals, e.g., submitted to the National Science Foundation (NSF), counting on the CESM's contributions to the CMIPs. Additionally, NCAR and CESM assist the community with the analysis of CMIP simulations by providing a CMIP Analysis Platform (https://www2.cisl.ucar.edu/resources/cmip-analysis-platform) as well as by making many diagnostics tools available for use of the community such as the Climate Variability and Diagnostics Package (CVDP). Of course, high level of scrutiny and analysis of the CESM simulations, in turn, feeds back to CESM, promoting further model development as well as enhancing collaborations.

As indicated in the Introduction, CESM2.0 has been released to the community in early June 2018. This version forms the code base for our CMIP6 simulations – that is, no code changes (i.e., no physics changes or tunings, parameter changes, etc.) are allowed with the possible exception of bug corrections. The CMIP6 pre-industrial simulations require spun-up land and ocean biogeochemisty (BGC) fields for initialization. These spun-up states were obtained in mid-July 2018, permitting us to start our pre-industrial control simulations. During the process of obtaining spun-up states for ocean BGC, some deficiencies in ocean BGC simulations were identified, resulting largely from deficiencies in ocean circulation. Tuning of ocean BGC was allowed to partially alleviate ocean BGC deficiencies. A new model version that incorporates these spun-up states along with additional non-answer-changing diagnostics and component sets relevant for many other CMIP6 simulations will be released in Fall 2018 as CESM2.1.

In following internationally established CMIP protocols, specific simulations (e.g., a long, 500-1000 years, pre-industrial control and at least 1850 – near present coupled simulation) are required for the scientific release of CESM2. Noting that these simulations are also part of the core simulations for CMIP6 (referred to as the DECK simulations, see below), 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 in the next section.

The CESM CMIP6 primary simulations consist of i) the DECK simulations – long pre-industrial control; 1%-increase in CO_2; instantaneous increase to $4xCO_2$; and a specified sea surface temperature (SST) simulation, 1979-present – and ii) the highest priority, i.e., Tier 1, experiments from many MIPs, each of which has been carefully reviewed by the WGs with the approval of the SSC (Tables 1 and 2). Note that the DECK simulations have to be performed for most of the configurations that will be part of CMIP6. By participating in CMIP6 (in the DECK and MIPs), 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.

As discussed below, there is much more emphasis on science in CMIP6 relative to previous CMIPs. The MIPs, following from the DECK, target a range of specific questions and aim to fill the scientific gaps of the previous CMIP phases. These additional simulations are usually listed under Tier 2 and Tier 3 simulations which are not included in our broader CESM CMIP6 Simulations request. Because of the interest of the broader CESM community in these science questions, many of the WGs plan to devote their computational resources to perform various Tier 2 and Tier 3 simulations as requested in their proposals.

Beyond the nominal 1° version of CESM2, experience will be gained from performing simulations at finer horizontal and vertical resolutions for the atmosphere, ocean, sea-ice, land-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, and topographically-forced circulations. In particular, CAM5 (Community Atmosphere Model version 5) simulations found that midlatitude winter and spring extreme precipitation over land are significantly better represented in the high-resolution model configuration (Wehner et al. 2014). 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. 2014). As articulated in the Supplementary Material, many WGs propose high resolution development and evaluation efforts towards a high-resolution CESM2 version.

The CESM WGs that are primarily responsible with the component models plan to continue their development efforts during this allocation cycle towards their inclusion in CESM3. Among these WGs, the OMWG is in the midst of a major transition of the dynamical core of the CESM ocean component, moving from the Parallel Ocean Program (POP) that has been used in CESM since CCSM version 2 through the present, to the MOM6. The choice of MOM6 as the new dynamical core was finalized shortly after the previous CSL proposal was submitted. During the later-half of the previous allocation cycle, work began in earnest in establishing a collaborative development arrangement with the MOM6 team at the Geophysical Fluid Dynamics Laboratory (GFDL), configuring a prototype version of MOM6 for CESM and developing the necessary infrastructure to couple MOM6 to the CESM framework. The efforts during this proposal period will include completing the porting of the physical parameterizations developed in POP to MOM6, exercising MOM6 in a variety of configurations within the CESM framework to establish a better understanding of its simulation biases and sensitivities, and to begin to exploit the new capabilities that are enabled by the MOM6 dynamical core. These activities will be performed in close collaboration with the sea-ice and land-ice modeling folks, involving the PCWG and LIWG. The development efforts towards CESM3 also include bringing Marine Biogeochemistry Library (MARBL) into MOM6; re-evaluation of coupling methods between sea-ice and the ocean model; and incorporating new capabilities in the Community Ice Sheet Model (CISM) with the goal of supporting climate simulations with a dynamic Antarctic ice sheet.

CMIP6 Simulations

Request for 147.0 M core-hours in Year 1

The CMIP6 experimental design can be found at http://www.wcrp-climate.org/wgcm-cmip/wgcm-cmip6 and is discussed in detail in Eyring et al. (2016). With the Grand Challenges of the World Climate Research Programme (WCRP) as its scientific backdrop, CMIP6 focuses on three broad science 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 these questions, CMIP6 consists of two major elements: i) the DECK and CMIP historical simulations (1850 – near present) that will maintain continuity and help document basic characteristics of models across different phases of CMIP; and ii) an ensemble of CMIP-Endorsed 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. Indeed, many of the MIPs involve simulations that are particularly designed to address specific science questions that arose from the previous MIP exercises. Based on the community input via the WGs and subsequent discussions and prioritizations in WGs and SSC, CESM will be participating in the CMIP6 MIPs listed in Table 1, using primarily the nominal 1° version of CESM2. Under the CESM2 CMIP6 Simulations labeled portion of the current proposal, only the DECK and MIPs Tier 1 (highest priority) simulations are covered. Any other Tier 2 and Tier 3 simulations that WG members (CESM community) are interested in performing are included in their respective requests.

MIP acronym	MIP name	Name of primary sponsor(s)
AerChemMIP	Aerosols and Chemistry Model Intercomparison Project	Lamarque/Emmons/Liu (Wyoming)
C4MIP	Coupled Climate Carbon Cycle Model Intercomparison Project	Lindsay
CDRMIP	Carbon Dioxide Removal Model Intercomparison Project	D. Lawrence/Lindsay
CFMIP	Cloud Feedback Model Intercomparison Project	Medeiros/Kay (CU)/Klein (LLNL)
DAMIP	Detection and Attribution Model Intercomparison Project	Tebaldi/Arblaster
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/Otto-Bliesner
LS3MIP	Land Surface, Snow and Soil Moisture	D. Lawrence
LUMIP	Land-Use Model Intercomparison Project	D. Lawrence/P. Lawrence
OMIP	Ocean Model Intercomparison Project	Danabasoglu/Lindsay
PAMIP	Polar Amplification Model Intercomparison Project	Deser/Philips
PMIP	Palaeoclimate Modelling Intercomparison Project	Otto-Bliesner
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 Stratosphere-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 MIPs for which CESM plans to perform their Tier 1 simulations. Only partial participation is planned for HighResMIP and CORDEX. The scientists who are primarily responsible for a given MIP are listed under the rightmost column. These *sponsors* are responsible for ensuring that i) simulations are correctly performed, ii) data are correctly posted, and iii) analysis is performed. Data only MIPs are only requesting specific output streams, but no additional simulations.

The complete collection of CESM2 simulations (DECK + MIPs Tier 1) is summarized in Table 2. It is important to note that the participation of CESM in 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 in CMIP6 is an extraordinary opportunity to document this new version, akin to the participation of previous CCSM and CESM versions in CMIPs. 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 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).

MIP name	Expt. Name	CMIP6 Exp Name	# years per ensemble member	#realizations	CESM2-BGC (1-degree)	CESM2 WACCM (1-degree)	CESM2 SE (1/4-degree)	B case	F case	G case	I case
DECK	Control	piControl	1000	1	1	0	0	1	0	0	0
	Control	esm-piControl	500	1	1	0	0	1	0	0	0
	Control WACCM	piControl	500	1	0	1	0	1	0	0	0
	Control high-res	piControl	175	1	0	0	1	1	0	0	0
	1% to 4x	1pctCO2	150	1	1	1	0	1	0	0	0
		1pctCO2									
	4xCO2	abrupt-4xCO2	150	1	1	1	0	1	0	0	0
		abrupt-4xCO2									
	AMIP (1979-2014)	amip	35	1	1	1	0	0	1	0	0
		amip									
Historical	1850-2014	historical	165	10	1	0	0	1	0	0	0
	1850-2014	esm-hist	165	3	1	0	0	1	0	0	0
Historical WACCM	1850-2014	historical	165	3	0	1	0	1	0	0	0
AerChemMIP	hist-piNTCF	hist-piNTCF	165	1	0	1	0	1	0	0	0
	hist-1950HC	hist-1950HC	65	1	0	1	0	1	0	0	0
	histSST	histSST	165	1	0	1	0	0	1	0	0
	histSST-piNTCF	histSST-piNTCF	165	1	0	1	0	0	1	0	0
	histSST-1950HC	histSST-1950HC	65	1	0	1	0	0	1	0	0
	histSST-piCH4	histSST-piCH4	165	1	0	1	0	0	1	0	0
	ssp370	ssp370	85	2	0	1	0	1	0	0	0
	ssp370-lowNTCF	ssp370-lowNTCF	85	3	0	1	0	1	0	0	0
	ssp370SST	ssp370SST	41	1	0	1	0	0	1	0	0
	ssp370SST-lowNTCF	ssp370SST-lowNTCF	41	1	0	1	0	0	1	0	0
	ssp370SST-lowAer	ssp370SST-lowAer	41	1	0	1	0	0	1	0	0
	ssp370SST-lowBC	ssp370SST-lowBC	41	1	0	1	0	0	1	0	0
	ssp370SST-lowO3	ssp370SST-lowO3	41	1	0	1	0	0	1	0	0
	piSSTclim	** piSST	30	1	0	1	0	0	1	0	0
	piSSTclim-NTCF	piClim-NTCF	30	1	0	1	0	0	1	0	0
C4MIP	1%BGC	1pctCO2-bgc	140	1	1	0	0	1	0	0	0
	esmssp5-85	esm-ssp585	85	1	1	0	0	1	0	0	0
CFMIP	amip-p4K	amip-p4K	36	1	1	0	0	0	1	0	0
	amip-4xCO2	amip-4xCO2	36	1	1	0	0	0	1	0	0
	amip-future4K	amip-future4K	36	1	1	0	0	0	1	0	0
	aqua-control	aqua-control	10	1	1	0	0	0	1	0	0
	aqua-p4K	aqua-p4K	10	1	1	0	0	0	1	0	0
	aqua-p4xCO2	aqua-4xCO2	10	1	1	0	0	0	1	0	0
DAMIP	histALL	historical-ext	6	10	1	0	0	1	0	0	0
	histNAT	hist-nat	165	3	1	0	0	1	0	0	0
		hist-nat									
	histGHG	hist-GHG	165	3	1	0	0	1	0	0	0
		hist-GHG									
	histAER	hist-aer	165	3	1	0	0	1	0	0	0
DCPP	A1	dcppA-hindcast	300	10	1	0	0	1	0	0	0
	C1.1	dcppC-atl-control	10	25	1	0	0	1	0	0	0
	C1.2	dcppC-amv-pos	10	25	1	0	0	1	0	0	0
	C1.3	dcppC-amv-neg	10	25	1	0	0	1	0	0	0
	C1.4	dcppC-pac-control	10	10	1	0	0	1	0	0	0
	C1.5	dcppC-ipv-pos	10	10	1	0	0	1	0	0	0
	C1.6	dcppC-ipv-neg	10	10	1	0	0	1	0	0	0
	C3.1	dcppC-hindcast-noPinatubo	10	10	1	0	0	1	0	0	0
	C3.4	dcppB-forecast	10	10	1	0	0	1	0	0	0

MIP name	Expt. Name	CMIP6 Exp Name	# years per ensemble member	#realizations	CESM2-BGC (1-degree)	CESM2 WACCM (1-degree)	CESM2 SE (1/4-degree)	B case	F case	G case	I case
GeoMIP6	G1ext	G1	100	1	0	1	0	1	0	0	0
	G6sulfur	G6sulfur	81	3	0	1	0	1	0	0	0
	G6solar	G6solar	81	3	0	1	0	1	0	0	0
GMMIP	AMIP20C (1870-2014)	amip-hist	145	1	0	0	1	0	1	0	0
HighResMIP	AMIP (1950-2014)	hist-1950	65	1	0	0	1	0	1	0	0
ISMIP6	PI Control	piControl-withism	300	1	1	0	0	1	0	0	0
	1% CO2	1pctCO2-withism	350	1	1	0	0	1	0	0	0
LS3MIP	Land-Hist	land-hist	165	1	1	0	0	0	0	0	1
	LFMIP-pdLC	lfmip-pdLC	121	1	1	0	0	1	0	0	0
	LFMIP-rmLC	lfmip-rmLC	121	1	1	0	0	1	0	0	0
LUMIP	deforest-glob	deforest-globe	80	1	1	0	0	1	0	0	0
	Historical no LULCC	hist-noLULCC	165	1	1	0	0	1	0	0	0
	SSP3-7 w/ SSP1-2.6 LU	ssp370-ssp126Lu	85	1	1	0	0	1	0	0	0
	SSP1-2.6 w/ SSP3-7 LU	ssp126-ssp370Lu	85	1	1	0	0	1	0	0	0
	esm SSP5-8.5 w/ SSP1-2.6 LU	esm-ssp585-ssp126Lu	85	1	1	0	0	1	0	0	0
	land-hist-alt-start-year	land-hist-alt-start-year	165	1	1	0	0	0	0	0	1
	Land-Hist w No LULCC	land-NoLu	165	1	1	0	0	0	0	0	1
OMIP	OMIP1	OMIP-A	310	1	1	0	0	0	0	1	0
PMIP4	Last Millenium	past1000	1000	3	1	0	0	1	0	0	0
	Mid Holocene	midHolocene	700	1	1	0	0	1	0	0	0
	Last Glacial Maxiimum	lgm	700	1	1	0	0	1	0	0	0
	Last Interglacial	lig127k	700	1	1	0	0	1	0	0	0
	Mid-Pliocene	midPliocene-eoi400	700	1	1	0	0	1	0	0	0
RFMIP	RFMIP-ERF-PI-Cntrl	piClim-control	30	1	1	0	0	0	1	0	0
	RFMIP-ERF-Anthro	piClim-anthro	30	1	1	0	0	0	1	0	0
	RFMIP-ERF-GHG	piClim-ghg	30	1	1	0	0	0	1	0	0
	RFMIP-ERF-AerO3	piClim-aer	30	1	1	0	0	0	1	0	0
	RFMIP-ERF-LU	piClim-lu	30	1	1	0	0	0	1	0	0
	RFMIP-ERF-4xCO2	piClim-4xCO2	30	1	1	0	0	0	1	0	0
ScenarioMIP	SSP5-8.5	ssp585	85	1	1	1	0	1	0	0	0
		ssp585									
	SSP3-7	ssp370	85	1	1	1	0	1	0	0	0
		ssp370									
	SSP2-4.5	ssp245	85	1	1	1	0	1	0	0	0
		ssp245									
	SSP1-2.6	ssp126	85	1	1	1	0	1	0	0	0
		ssp126									
VolMIP	volc-long-eq	volc-long-eq	20	9	1	0	0	1	0	0	0
	volc-pinatubo-full	volc-pinatubo-full	3	25	1	0	0	1	0	0	0
	volc-pinatubo-surf	volc-pinatubo-surf	3	25	1	0	0	1	0	0	0
	volc-pinatubo-strat	volc-pinatubo-strat	3	25	1	0	0	1	0	0	0
PAMIP	pa-pdSSTpdSIC	pa-pdSSTpdSIC	1	100	1	0	0	0	1	0	0
	pa-piSSTpdSIC	pa-piSSTpdSIC	1	100	1	0	0	0	1	0	0
	pa-pdSSTpiSICArctic	pa-pdSSTpiSICArctic	1	100	1	0	0	0	1	0	0
	pa-pdSSTpiSICAntarctic	pa-pdSSTpiSICAntarctic	1	100	1	0	0	0	1	0	0
	pa-pdSSTfutSICArctic	pa-pdSSTfutSICArctic	1	100	1	0	0	0	1	0	0
	pa-pdSSTfutSICAntarctic	pa-pdSSTfutSICAntarctic	1	100	1	0	0	0	1	0	0
CDRMIP	4x to 1x CO2	1pctCO2-cdr	200	1	1	0	0	1	0	0	0
	Instantaneous CO2 removal	esm-pi-cdr-pulse	100	1	1	0	0	1	0	0	0
	Instantaneous CO2 addition	esm-pi-co2-pulse	100	1	1	0	0	1	0	0	0

Timings (core-hours/yr)	CESM2-CAM6-BGC 1deg B-case	3500			
	CESM2-CAM6-BGC 1deg F-case	3000			
	CESM2-WACCM-BGC 1deg B-case	28000			
	CESM2-WACCM-BGC 1deg F-case	25000			
	CESM2 SE 1/4deg B-case	150000			
	CESM2 SE 1/4deg F-case	120000			
	OCN G-case	1000			
	CLM I-case	1000			
		B-case	F-case	G-case	I-case
Total # sim.years	CESM2-CAM6-BGC	18302	953	310	495
	CESM2-WACCM-BGC	2876	860		
	CESM2 SE 1/4deg	175	210		
Total cost (M pe-hours)	CESM2-CAM6-BGC	67.72			
	CESM2-WACCM-BGC	102.03			
	Total 1-degree models	169.75			
	CESM2 SE 1/4deg	51.45			

Table 2. CESM2 DECK and MIPs Tier 1 simulations. The bottom panel includes timings and the total number of simulations for each configuration as well as the total time needed in M of core-hours.

In addition to the CESM2 standard 1° model version, we will use CESM2-WACCM6 (Whole Atmosphere Community Climate Model version 6) with enhanced middle and upper atmosphere representation (at 1° atmosphere and ocean horizontal resolutions) to simulate a variety of fields (ozone, volcanic aerosols, and 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. Therefore, computational resources are requested for those simulations as well.

Overall, the CMIP6 simulations with the nominal 1° version of the model will require about 169.8 M core-hours (Table 3). However, with the finalization of the CESM2 and its readiness for CMIP6 in late July 2018, we have started some of the DECK simulations already, using our existing allocation. Specifically, CESM pre-industrial control simulations both with WACCM6 and CAM6 atmospheric models and three ensemble members of the historical simulation with WACCM6 are progressing. With our remaining allocation until 31 October 2018, we anticipate nearly completing these simulations. We also anticipate completing several ensemble members of the historical simulation with CAM6 during this period. Thus, by the time the new allocation cycle starts, we believe that the total amount of core-hours needed for this purpose will be down to about 147.0 M.

These CESM2 CMIP6 simulations represent about 1/3 of our two-year CSL request. However, this request does not include possible contributions by CESM2 to the High-Resolution Model Intercomparison Project (HighResMIP). The primary reason is that the delay in finalizing CESM2 has prevented dedicated work on the high-resolution version of the model both from the individual component models and fully-coupled model perspectives. The WG requests do include various development and evaluation projects towards finalizing a CESM2 high-resolution version. As shown in Table 2, CESM2's participation in HighResMIP is anticipated to require about 51.5 M core-hours – based on Cheyenne estimates. We plan to seek other computational resources for these simulations.

The estimated storage need for CMIP6 DECK and MIPs Tier 1 simulations is about 5.7 PB. The spread sheet detailing these estimates are available at http://www.cesm.ucar.edu/CMIP6/data_estimate. Additional information is provided in the Data Management and Archive Requirements section below.

Community Projects

Request for 43.3 M core-hours in Year 2

Over the last three cycles of CSL proposals, we have defined a collection of experiments as *Community Projects* because they represent large simulations that are of interest to multiple WGs, and thus, to the broader CESM community, including international partners / collaborators. These community projects compliment the CMIP6-related community efforts, pushing community-wide scientific objectives further. Following the same philosophy, we have defined 5 projects for the current request, ranging from 2.8 to 15 M core-hours (see Table 3). The process for selection consisted of a call for proposals; a review of feasibility and cross-WG interest by the WG co-chairs; and finally a review and selection / recommendation by the SSC.

The following is a list of the selected topics. Further details of each community project request are available in the Supplementary Material document.

C1. <u>Transient Holocene (11.4 M core-hours):</u> The transient Holocene simulation will provide model data to the CESM community to more fully explore multidecadal and longer variability and rapid transitions of, for example: ENSO and other modes of climate variability; monsoons and droughts; the Atlantic meridional overturning circulation; and tropical/extratropical linkages. This unprecedented simulation, covering the period from 9000 years ago until the beginning of the last millennium simulation (850 AD) – for a total of 8150 years – will be run with the CESM2.1 FV2x1 model version. This simulation was included as a community project simulation in the last CESM CSL allocation request, but needed to be deferred due to the delay in CESM2. Requested by PaleoWG and LIWG with AMWG and WAWG interests in the behavior of the FV2 model version.

C2. <u>High-resolution ocean (POP) with biogeochemistry (BGC) (15.0 M core-hours):</u> Performing a hindcast eddy-resolving (0.1°) ocean simulation with BGC remains a high-priority. This is a frontier modeling capability and only a few such integrations have been conducted worldwide. A broad community of researchers has interest in these integrations. The model will use MARBL and will be forced with the newly developed JRA55 inter-annually varying atmospheric data sets based on the Japanese Meteorological Agency Reanalysis Product. The simulation will be performed for the 50-year period from 1967 – 2016. Requested by OMWG, PCWG, and BGCWG; directly feeds in to the world-wide ocean modeling community and the activities / efforts of the International CLIVAR Ocean Model Development Panel (OMDP).

C3. <u>Subseasonal-to-seasonal (S2S) hindcasts (2.8 M core-hours):</u> S2S hindcasts for years 1999 - 2018 will be carried out using the CESM2 nominal 1° version and following the Subseasonal Experiment (SubX) protocol. 10-member ensemble hindcasts will be initiated every Wednesday from March to mid-October and run for 45 days each. Comparison of hindcasts with CESM2 to CESM1 will allow for the characterization in changes in model skill coming strictly from the new model physics in CESM2. The analysis will focus on simulation skill of weeks 3 to 6 surface temperature and precipitation as well as the skill of the Madden-Julian Oscillation, which is improved in climate simulations with CESM2 as compared to CESM1. Requested by AMWG with interests from OMWG and PCWG; directly complimenting the national SubX efforts.

C4. <u>CESM2 with RCP8.5 projections (6.3 M core-hours):</u> CESM2 has considerable changes relative to the CESM1. In order to understand these changes and their impact on climate variability and change, simulations are needed with identical forcing for the

20th and 21st centuries. Given this, an ensemble set of historical and future projections that apply the same forcing as in the CESM-LE configuration is proposed. Specifically, a 10-member ensemble of CESM2 simulations for the 1920-2100 period, forced by the same historical forcing as the CESM-LE and the RCP8.5 projection scenario will be performed, allowing comparisons of the historical and projected climate change under RCP8.5 with the projected climate change in the standard CMIP6 runs with the CESM2 using SSP5-8.5 and SSP3-7.0. Given the huge success and wide use of the CESM-LE project, having the opportunity to establish how the CESM2 projections differ from the CESM-LE projections will be very beneficial to a large community of scientists. Requested by PCWG with significant interests from CVCWG and AMWG.

C5. Development of a CESM Arctic Prediction System (CAPS) (7.8 M core-hours): CESM is uniquely placed to advance understanding processes and predictability of polar regions. Predictability on scales from sub-seasonal to seasonal (S2S) out to decadal and centennial requires a high resolution coupled system. CESM2 with a refined mesh atmosphere at high resolution (7-15 km) can be coupled to high-resolution versions of the other component models. Before doing so, however, we propose to do initial tests and development with a 1° ocean with BGC coupled to a high-resolution atmosphere. As such, the request represents the start of a larger development of a coupled prediction system for the Arctic that would involve high resolution prediction with initialized forecasts when a more unified system is ready. It is consistent with one of NSF's 10 big ideas, i.e., Navigating the New Arctic. Requested by AMWG, LIWG, LMWG, and PCWG.

Working Group Research Objectives and Requests

Request for 83.6 M core-hours in Year 1 and 186.8 M core-hours in Year 2

In this section, we summarize the overall research goals and objectives specific to each WG. In addition, for each WG, we provide the requested computing allocation, split between development and production as well as between Years 1 and 2 – identified as D-Y1, D-Y2, P-Y1, and P-Y2, respectively). All the core-hour requests listed in Table 3 are in thousands of Cheyenne core-hours. Further details of each WG's request are available in the Supplementary Material document.

	Year 1		Year 2		Total per WG		
	Dev	Prod	Dev	Prod	Year 1	Year 2	Y1+Y2
CMIP6 DECK and MIP Tier 1 Exp		147000			147000		147000
Working group							
AMWG	8100	3320	24700	2240	11420	26940	38360
BGCWG	2710	5105	10590	7740	7815	18330	26145
CHWG	1860	1650	4000	4760	3510	8760	12270
CVCWG	0	11000	0	24421	11000	24421	35421
LIWG	750	4154	4514	6500	4904	11014	15918
LMWG	1673	5324	5070	11303	6997	16373	23370
OMWG	2570	4497	12461	3592	7067	16053	23120
PaleoWG	3500	6300	11250	12250	9800	23500	33300
PCWG	2218	2097	4046	6153	4315	10199	14514
SEWG	6000	0	6000	0	6000	6000	12000
WAWG	5223	5540	10969	14216	10763	25185	35948
Total					230591	186775	417366
Community Projects							
C1 Transient Holocene				11400		11400	11400
C2 High-res Ocean with BGC				15000		15000	15000
C3 S2S Hindcasts				2800		2800	2800
C4 CESM2 with RCP8.5 Projections				6300		6300	6300
C5 Dev. Arctic Prediction System				7800		7800	7800
Total				43300		43300	43300
Total (CMIP6 + dev + prod + comm)					230591	230075	460666
Target					230000	230000	460000

Table 3. Complete list of core-hour allocation requests for CMIP6 DECK and MIPs Tier 1; WG development and production; and community projects. The entries are in thousands of Cheyenne core-hours. The requests for Year 1 and Year 2 as well as the total for both years are provided.

For all requests, the choice was left to the WGs to balance between simulation throughput and cost (see Model Performance section below). Consequently, slightly different estimates can be found for the same model configuration. In addition, the estimates for new versions have relied on simple scaling arguments, considering changes, for example, in number of levels, tracers, or horizontal resolution.

Atmosphere Model Working Group (AMWG)

D-Y1: **8.1M;** D-Y2: **24.7M;** P-Y1: **3.3M;** P-Y2: **2.2M; Total: 38.3M**

The AMWG utilizes CSL resources primarily for the development of the 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 also covers sensitivity experiments aimed at understanding the many interactions among the represented physical and dynamical processes across climate regimes and multiple timescales. AMWG development activities will include: i) Analyzing impacts of CAM6 physics schemes on current model behavior;

ii) Refocusing on model resolution – both vertical and horizontal; iii) Evaluating new candidate atmospheric dynamical cores - MPAS (Model for Prediction Across Scales) and FV3 (Finite Volume Cubed-Sphere); iv) Developing scientifically-supported regionally-refined model configurations for domains of interest such as the Continental U.S. (CONUS) and the west Pacific warm-pool; v) Continuing physics developments including improvements to Cloud Layers Unified By Binormals (CLUBB) as well as exploring alternative convection, turbulence, and drag schemes; and vi) Exploring new frameworks for climate model evaluation including idealized configurations such as radiative-convective equilibrium and gray-radiation physics, cloud-locking capabilities, forecast configurations, nudging, and coupling with the Data Assimilation Research Testbed (DART) framework. On the production side, AMWG plans to perform simulations to reconstruct the evolution of CAM6 from its previous version (CAM5) to understand the behavior of the new model and participate in Tier 2 of the Cloud Feedback Model Intercomparison Project (CFMIP).

Biogeochemistry Working Group (BGCWG)

D-Y1: **2.7M;** D-Y2: **10.6M;** P-Y1: **5.1M;** P-Y2: **7.7M; Total: 26.1M**

The goal of the BGCWG is to produce a state-of-the-art Earth system model for the research community that includes terrestrial and marine ecosystem biogeochemistry. This model will be used to explore ecosystem and biogeochemical dynamics and feedbacks in the Earth system under past, present, and future climates. Land and ocean ecosystems influence climate through a variety of biogeophysical and biogeochemical pathways. Interactions between climate and ecosystem processes, especially in response to human modification of ecosystems and atmospheric CO_2 growth, produce a rich array of climate forcings and feedbacks that amplify or diminish climate change. At present only about half of anthropogenic carbon remains in the atmosphere to drive climate change; the remainder is removed in about equal amounts by the land biosphere and the oceans. While the magnitude of contemporary ocean uptake of anthropogenic carbon is constrained by observations to within 10%, the future uptake is uncertain. Thus, a primary objective of the BGCWG 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 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. Therefore, another primary objective of the BGCWG is to analyze these, and other, terrestrial feedbacks using CESM. Subsequently, better understanding of ecosystem and biogeochemical dynamics and feedbacks with respect to a changing climate requires an expansion of current CESM land and ocean model capabilities. Thus, biogeochemistry development is focused on continued development of the Newton-Krylov fast spin-up technique; preliminary high-resolution experiments; continued development of biogeochemical parameterizations; porting of MARBL to the new ocean component MOM6; coupling across components and understanding interactions; and automated techniques for the optimization of model parameters. Production runs address fully coupled carbon cycle experiments and single component experiments with well-established models. BGCWG is also requesting computing resources to perform several Tier 2 simulations for MIPs of interest and for additional carbon cycle sensitivity experiments.

Chemistry Climate Working Group (CHWG)

D-Y1: **1.9M;** D-Y2: **4.0M;** P-Y1: **1.7M;** P-Y2: **4.8M; Total: 12.4M**

The goal of the CHWG 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, and to understand the role of climate change on tropospheric composition. The development and production simulations requested here will lead to improving the representation of tropospheric composition and air quality. Inorganic nitrate aerosols have been 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. Increasing horizontal resolution is a key factor in improving air quality simulations, so an important component of CESM chemistry development is the testing and tuning of the CAM-SE (spectral element) configuration with regional refinement and comprehensive chemistry. In addition, the CHWG will continue with testing of CSLAM (Conservative semi-Lagrangian Multi-tracer) in CAM-Chem and the impact of this transport scheme on tracer transport, chemistry, and aerosols. Biogenic emissions will also be updated in MEGAN (Model of Emissions of Gases and Aerosols from Nature). Simulations with CAM-Chem will be evaluated against field experiments, and performing analyses of the responses to weather and climate changes. The capability for running daily to near-seasonal air quality forecasts will be developed in CAM-Chem. CHWG will finalize the development of CAM-Chem with the CESM and DART multi-instances framework, including the testing of the best configuration in terms of the number of nodes per instances, which has not been tested for chemistry yet. The goal is to assimilate satellite observations from the day before using CAM-Chem – DART to initialize a deterministic forecast of several days. Updates to aerosol dry deposition and representation of brown carbon will also be tested in CAM6-Chem. In addition, simulations with the expanded very short-lived (VSL) halogen chemistry will be performed for comparison to CMIP6 simulations. Operational-style short-term and seasonal forecasts of air quality will be run during both years of the request.

Climate Variability and Change Working Group (CVCWG)

D-Y1: **0M**; D-Y2: **0M**, P-Y1: **10.5M;** P-Y2: **24.9M; Total: 35.4M**

The goals of the CVCWG 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 runs 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, AMIP and *Pacemaker* style historical ensembles, and simulations for investigation into climate variability and extremes. Some specific simulation examples include: additional ensemble members for Tier 1 of ScenarioMIP to reduce uncertainty; one ensemble member each for Tier 2 and 3 DAMIP (Detection and Attribution Model Intercomparison Project) simulations; and additional runs for CFMIP with aqua planet simulations to dig deeper into feedbacks. The CVCWG has planned high-resolution simulations for Year 2 of this allocation. Also, the WG will perform long pre-industrial control simulations and large ensembles of historical simulations with a hierarchy of model configurations to explore and understand internally-generated patterns, time scales and mechanisms of climate variability and change. Another ongoing research area of the CVCWG is to gain an understanding of the interplay between external forcing and internal variability, with a focus on the Interdecadal Pacific Oscillation. The CVCWG is also investigating processes and mechanisms that characterize high impact events and how such events and driving mechanisms might change in the future. These events are of great significance to society with potential consequences to habitats, economics, and human life. Towards this effort, there are specific simulations planned to investigate tropical cyclones, precipitation and drought extremes, atmospheric rivers, and sea level rise.

Land Ice Working Group (LIWG)

D-Y1: **0.8M;** D-Y2: **4.5M;** P-Y1: **4.2M;** P-Y2: **6.5M; Total: 16.0M**

The primary objective of the LIWG in the timeframe of this proposal is to carry out pioneering simulations using the coupled ice-sheet – climate model consisting of the CISM version 2.1 (CISM2.1) in CESM2. This coupling will enable the LIWG to carry out the Tier 1 and Tier 2 coupled climate experiments of the Ice Sheet Model Intercomparison Project for CMIP6 (ISMIP6). ISMIP6 is an international effort whose goals are to estimate past and future sea level contributions from the Greenland and Antarctic ice sheets, along with associated uncertainty, and to investigate feedbacks due to dynamic coupling between ice sheet and climate models. The LIWG aims to run a full suite of stand-alone ice sheet and coupled ice sheet–climate experiments as specified in the ISMIP6 protocols. Production resources will also be used for coupled CISM-CESM runs on time scales of a millennium or longer: for paleoclimate simulations of the last deglaciation and for multi-century future simulations of Greenland deglaciation. Running at coarse atmosphere resolution (2°) with accelerated ice sheet dynamics will allow multi-millennial simulations. Meanwhile, in-house development of new science capabilities will continue, with the focus expanding from the Greenland ice sheet to include Antarctic and paleo ice sheets. These efforts will broaden to include coupled simulations on a variable-resolution grid – where the horizontal resolution of CESM is enhanced over ice sheets – to better capture the surface mass balance (SMB) of Greenland and West Antarctica. Higher resolution improves ice sheet surface climate and precipitation and, in the long term, will render the use of elevation classes over ice sheets unnecessary when coupling CESM to CISM. Other development runs will test new CISM capabilities to simulate marine ice sheets, with the long-term goal of supporting climate simulations with a dynamic Antarctic ice sheet in CESM3.

Land Model Working Group (LMWG)

D-Y1: **2.0M;** D-Y2: **5.1M;** P-Y1: **7.0M;** P-Y2: **11.0M; Total: 25.0M**

The goals of the LMWG are 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 LMWG request. The LMWG has pursued an ambitious program of model development, which culminated with the release of the Community Land Model version 5 (CLM5) in February 2018. Several additional large development projects have been progressing in parallel to CLM5 development including a multi-layer canopy scheme, a representative hillslope hydrology model, and the Functionally Assembled Terrestrial Ecosystem Simulator (FATES) configuration of CLM. These projects will continue into the next CSL allocation cycle along with other development projects on water management and agriculture model development. Parameter estimation/calibration is an increasingly important feature of CLM development. In addition, the LMWG in collaboration with land modeling scientists across NCAR has begun work towards unifying land modeling activities across NCAR to form the Community Terrestrial Systems Model (of which CLM5 is the current climate configuration of the broader CTSM). On the production side, the LMWG will participate in several MIPs as 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). CESM2 will participate in each of these MIPs, utilizing CLM5 in both coupled and land-only experiments.

Ocean Model Working Group (OMWG)

D-Y1: **2.6M;** D-Y2: **12.5M;** P-Y1: **4.5M;** P-Y2: **3.6M; Total: 23.2M**

The primary goals of the OMWG are to advance the capability and fidelity of the CESM ocean component in support of specific science objectives of the broader CESM community and to conduct curiosity driven research using CESM to improve our understanding of ocean processes, the role of the ocean in the Earth system, and its interactions with other Earth system components. The OMWG is in the midst of a major transition of the dynamical core of the CESM ocean component, moving from POP to MOM6. The choice of MOM6 as the new dynamical core was finalized shortly after the previous CSL proposal was submitted. During the later-half of the previous allocation cycle, work began in earnest in establishing a collaborative development arrangement with the MOM6 team at the GFDL, configuring a prototype version of MOM6 for CESM and developing the necessary infrastructure to couple MOM6 to the CESM framework. As of August 2018, this ground work is very near completion. A large fraction of the resources requested for the coming proposal period will be dedicated to completing the porting of the physical parameterizations developed in POP to MOM6, exercising MOM6 in a variety of configurations within the CESM framework to establish a better understanding of its simulation biases and sensitivities, and to begin to exploit the new capabilities that are enabled by the MOM6 dynamical core. These include, for example, prognostic sea level change, higher vertical resolution near the sea surface, more accurate treatment of topography, and more flexible time stepping. New and ongoing parameterization development efforts are underway on several fronts and will be transitioned to the MOM6 framework. Additionally, the transition to MOM6 opens the opportunity to broaden the use of CESM within the oceanographic research community and to begin to provide capabilities to the CESM community that had been difficult or impossible with the POP based dynamical core including idealized geometry configurations, regional or nested configurations, and configurations at eddy permitting resolutions. An objective of the development work proposed here is to be able to release one or more versions of CESM with MOM6 as the ocean component prior to the release of CESM3. Efforts using the DART with both moderate and high-resolution versions of CESM-POP will continue as a part of research in seasonal- to decadal prediction underway with CESM2. While no further development of the POP ocean component is planned, the resources requested for production experiments will exploit the newly available POP-based CESM2. A set of experiments will investigate the origin of bi-stability of the CESM2 climate expressed as a proclivity for shut down of Labrador Sea convection and expansion of sea ice over the subpolar North Atlantic. A separate line of research will investigate the physical basis for the very weak ventilation of the deep Pacific Ocean that became a critical problem in the representation of the marine carbon cycle in CESM2. A high-resolution coupled ocean-sea-ice simulation with POP will be completed using the recently prepared ocean forcing data set based on the JRA-55 (Japanese Reanalysis) product with sampling sufficient to investigate eddy-mean-flow interaction. An extension of work begun in the previous allocation using CESM1 for investigations of Atlantic meridional overturning circulation is also requested.

Paleoclimate Working Group (PaleoWG)

D-Y1: **3.5M;** D-Y2: **11.3M;** P-Y1: **6.3M;** P-Y2: **12.3M; Total: 33.4M**

The PaleoWG is a consortium of scientists interested in modeling and understanding past Earth climates, allowing a long-term perspective on Earth system feedbacks. PaleoWG members include participants from universities and laboratories, with interests ranging from early Earth to the recent millennium. The WG conducts scientific modeling experiments to establish relationships between forcings and feedbacks for specific time periods, and to explore the transient nature of these responses. Comparing model results to observational data is an important component of the PaleoWG efforts. The PaleoWG development goal 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 WG will continue to develop and explore model parameterizations and capabilities to shed light on unanswered questions about past climates, and for out-of-sample testing and evaluation of the model parameterizations that are being used in projections of the future. In this proposal, the WG includes testing new configurations of CESM, such as the capability to simulate the inception and retreat of Greenland, North American, and Eurasian ice sheets with CISM2 coupled to CESM2, the sensitivity of the Last Millennium simulation to alternate forcings, the application of CESM2 to the high CO_2 early Eocene, and the extension of the PMIP4 (Paleoclimate Modeling Intercomparison Project) simulation to include the water isotope implementation scheduled for CESM2.2. The PaleoWG production goal is to provide benchmark simulations of past climates to the community. These simulations offer the opportunity to test CESM2 for various forcing conditions, carry out detection and attribution studies, and improve confidence in its application for the future. The WG will carry out experiments as part of international intercomparison projects – PMIP4 and ISMIP6. The proposed production simulations are the Tier 2 and 3 simulations, 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. The PaleoWG simulations will utilize the FV 1°x1° version of CESM2 for several reasons. First, the protocols for PMIP and DeepMIP require the CMIP6 pre-industrial and 4xCO_2 simulations to be run by the modeling groups. These will be already completed as part of the CMIP6 DECK simulations being run at NCAR. Second, a finer ocean resolution of 1° has been shown to be necessary to simulate flow through narrow gateways. Third, the only low-resolution version of CESM2 will be the FV 2°x1° version, which is still under development.

Polar Climate Working Group (PCWG)

D-Y1: **2.2M;** D-Y2: **4.0M;** P-Y1: **2.1M;** P-Y2: **6.2M; Total: 14.5M**

The PCWG is a consortium of scientists who are interested in understanding and modeling Arctic and Antarctic climate and its relationship to global climate. To enable polar science within the PCWG and the CESM project as a whole, we request computing resources for both polar-specific CESM parameterization development and polar-specific CESM scientific research. The overall development objective is to ensure that CESM has state-of-the-art abilities to simulate polar climate. The proposed development work includes analysis of a land-fast sea ice parameterization, development of improved treatments of snow on sea ice, analysis of biogeochemistry within sea ice, testing of new sea ice dynamics formulations, improvements based on diagnostic evaluation of CMIP6 models, and development of coupling between sea ice and the ocean model for CESM3. In collaboration with the AMWG, the PCWG proposes to evaluate and improve atmospheric model moist physics parameterizations with a specific emphasis on polar boundary layer, turbulence, clouds, and precipitation processes, using a CMIP-endorsed instrument simulator package that enables scale-aware and definition-aware comparisons between models and observations. Another development work concerns investigating the impacts of spectrally resolving the surface emissivity. The overarching PCWG production goal is to enable important and topical polar science research using CESM. This includes experiments of value to a large number of researchers that are related to polar prediction, integrating models and observations to enhance process understanding, and understanding coupled system interactions and feedbacks. The proposed experiments make use of both the CESM-LE and CESM2 configurations. This will allow for the diagnoses of important climate processes relative to the large number of simulations available for the CESM-LE and also enhanced understanding of new interactions within CESM2.

Software Engineering Working Group (SEWG)

D-Y1: **6.0M;** D-Y2: **6.0M;** P-Y1: **0M;** P-Y2: **0M; Total: 12.0M**

The role of the SEWG is 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 new CESM revision on all 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. Computing time is requested to carry out this important function throughout the various CESM versions that will be generated.

Whole Atmosphere Working Group (WAWG)

D-Y1: **5.2M;** D-Y2: **11.0M;** P-Y1: **5.5M;** P-Y2: **14.2M; Total: 35.9M**

The WAWG research plan involves development designed to continue the move towards a unified sun-to-earth modeling, i.e., WACCM, 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 bringing 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 will contribute to ISA-MIP (Interactive Stratospheric Aerosol Model Intercomparison Project) for comparison of models with interactive stratospheric aerosols, GeoMIP (Geoengineering Model Intercomparison Project) for geoengineering studies, and QBOi for comparison of models with interactive quasi-biennial oscillations. Production will also include WACCM simulations of the Last Millennium (850-1850) with interactive volcanic aerosols derived from emissions, and studies of space climate with WACCM-X.

Data Management and Archive Requirements

As part of this CSL proposal, each WG generated estimates of the data volume associated with each proposed development and production experiment (listed in the Supplementary Material). These estimates are summarized in Table 4 in Terabytes (TB). For a few experiments, it is possible that the WG storage estimates may differ for similar simulations due to differences in their output fields. As in the previous CSL proposals, we follow the 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 in to 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. As shown in Table 4, we anticipate generating 10.1 PB of data during the upcoming CSL cycle, inclusive of all CMIP6, WG development and production, and community projects. The spread sheet detailing the storage estimates for the CMIP6 DECK and MIPs Tier 1 simulations are available at http://www.cesm.ucar.edu/CMIP6/data_estimate.

	Year 1		Year 2		Total per WG		
	Dev	Prod	Dev	Prod	Year 1	Year 2	Y1+Y2
CMIP6 DECK and MIP Tier 1 Exp		5700.0			5700.0		5700.0
Working group							
AMWG	47.5	13.8	93.3	7.0	61.3	100.3	161.6
BGCWG	92.8	192.0	358.4	164.8	284.8	523.2	808.0
CHWG	27.0	53.0	70.0	63.0	80.0	133.0	213.0
CVCWG	0.0	134.1	0.0	733.1	134.1	733.1	867.2
LIWG	45.5	63.1	18.1	273.0	108.6	291.1	399.7
LMWG	1.5	63.1	4.9	43.7	64.6	48.6	113.2
OMWG	38.6	155.6	229.4	71.2	194.2	300.6	494.8
PaleoWG	54.5	87.2	167.2	190.8	141.7	358.0	499.7
PCWG	46.5	42.1	83.8	164.1	88.6	247.9	336.5
SEWG	0.0	0.0	0.0	0.0	0.0	0.0	0.0
WAWG	51.5	50.0	19.8	28.2	101.5	48.0	149.5
Total					6959.4	2783.8	9743.2
Community Projects							
C1 Transient Holocene				137.0		137.0	137.0
C2 High-res Ocean with BGC				110.0		110.0	110.0
C3 S2S Hindcasts				1.0		1.0	1.0
C4 CESM2 with RCP8.5 Projections				130.0		130.0	130.0
C5 Dev. Arctic Prediction System				10.0		10.0	10.0
Total				388.0		388.0	388.0
Total (CMIP6 + dev + prod + comm)					6959.4	3171.8	10131.2

Table 4. Complete list of storage estimates for CMIP6 DECK and MIPs Tier 1; WG development and production; and community projects. The entries are in TB. The estimates for Year 1 and Year 2 as well as the total for both years are provided.

a. Data archiving

Of the 10.1 PB of the total data volume expected to be generated from all proposed experiments, about 5.7 PB is associated with the CMIP6 DECK and MIPs Tier 1, and another 0.6 PB with MIPs Tier 2 and Tier 3 simulations. To manage this data volume, we will take advantage of the new infrastructure (primarily campaign storage) and modify our archival strategy:

Development: Output data will be primarily stored on the requested glade partition (see below). If necessary, output will be copied to campaign store. We expect that this will account for approximately 10-20% of the development output. In addition, the data will be removed 36 months after creation, unless retention is requested from the relevant WG co-chairs. *One-off* development experiments will be removed more quickly at the PIs and / or WG co-chairs' discretion.

Production: Output data will be stored on glade scratch space (model *history* data) and will be converted to *timeseries* format, using lossless compression, and these timeseries data will be archived to campaign storage, for a period of no more than five years. As these data get close to their time limit on campaign storage, key data will be archived to HPSS if deemed appropriate and necessary. These data will then be gradually cut back to 50% of their 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 WG 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 https://csegweb.cgd.ucar.edu/expdb2.0/cgi-bin/expList.cgi). This database 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 campaign storage and HPSS according to the CESM Data Management and Data Distribution Plan. The CISL SAM (Systems Accounting Manager) website will also be used to assist in managing CESM data.

As of this proposal writing, the campaign storage is just being rolled out. As indicated above, although it is expected that only key output data will be written onto HPSS, the majority of the CMIP6-related data sets will likely be archived on HPSS at the end of their time limit on the campaign storage. However, it is anticipated that this HPSS need will be much less than the 6.3 PB (5.7 PB + 0.6 PB for Tier 2 and Tier 3) produced by the CMIP6-related simulations.

b. Data distribution

Development: In general, output data will be made available only to the WG members that are directly involved with the experiments. For WG members that do not have access to CSL resources, these data will be made available via the Climate Data Gateway (CDG), formerly known as the Earth System Grid (ESG), or the NCAR Data Sharing Service, as appropriate.

Production: Output data will be made available according to the guidelines established by the CESM Data Management and Data Distribution Plan, which was formulated by the CESM SSC, NCAR, and NSF. Initially, access is restricted to the WG members directly involved with the experiments. After a period of no more than 12 months following creation, these data will be made available to the community via the CDG. The CMIP6 model output will be made available to the NSF-funded CISL CMIP Analysis Platform (https://www2.cisl.ucar.edu/resources/cmip-analysis-platform) and the Earth System Grid Federation (ESGF) in accordance with the CMIP protocol and requirements.

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 WG members who have CSL access 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 WG members and involvement.

d. GLADE project file space (total request: 1.2 PB)

In order to minimize the usage of HPSS for storing development results, we request 1.2 PB of CESM GLADE project space. This estimate is based on our previous experience and our anticipated use. 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 space is collectively managed by the CESM WGs.

e. Campaign storage space (total request: 6.6 PB = 2.6 PB + 4 PB for CMIP6)

In order to minimize the usage of HPSS for storing CESM output, we request 2.6 PB of CESM campaign storage space. This space will be used instead of HPSS to avoid unnecessarily writing files to tape. We also request 4 PB of campaign storage for CMIP6 output specifically. To help CISL partially cover the costs associated with the CMIP6-related storage, the CESM program requested and obtained funds from the NCAR directorate for CISL to purchase 4 PB of campaign storage. Developments in CDG that allow data distribution via campaign storage instead of HPSS will be used as much as possible.

Model Performance

The vast majority of the proposed simulations in this request use the CESM2 nominal 1° horizontal resolution version with CAM6. Therefore, it is critical to have a comprehensive understanding of the overall performance (cost and throughput) for this *work-horse* configuration. Primarily by optimizing the layout of each component and also by improving vectorization and reducing communication costs, we were able to achieve a very respectable rate of approximately 30 model years per wall-clock day on 4320 cores on Cheyenne (Fig. 1). Such throughput is absolutely necessary to enable the timely completion of all the proposed simulations, especially so for the CMIP6-related experiments. This, obviously, comes with an increased cost (about 3.5 K / year for 4320 tasks at 30 model years per wall-clock day vs. 2.2 K / year for 720 tasks at 8 model years per wall-clock day).

A similar approach was followed for the CESM2 nominal 1° horizontal resolution version with WACCM6. A rate of 6 model years per wall-clock day was achieved on 7056 processors, resulting in 28 K core-hours per year. As in the CAM6 version, this represents an increased cost compared to, for example, getting 4 years per day on 3528 processors at a cost of about 21 K core-hours per year. To re-iterate, increased throughput is crucial to complete the proposed simulations on time.

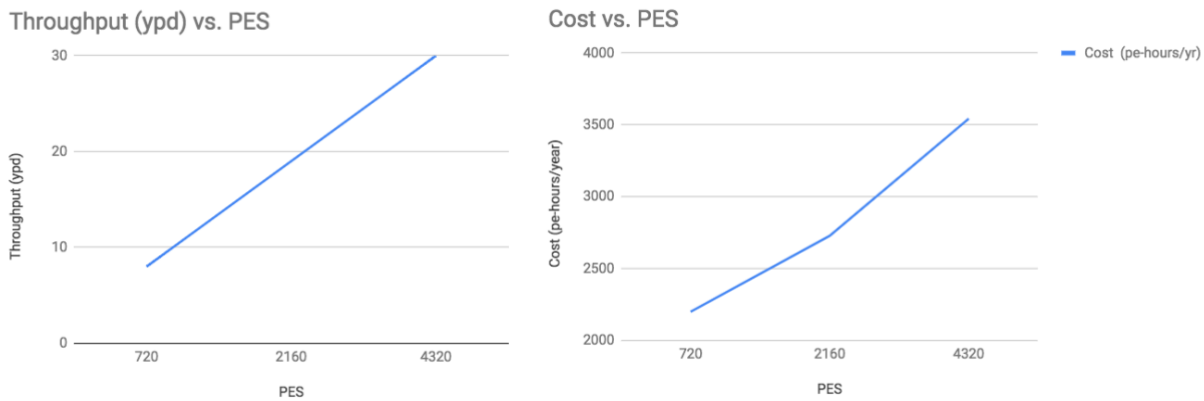

Figure 1. CESM2 nominal 1° version throughput (simulation years per wall-clock day) and cost (core-hours per simulation year) as a function of processor count.

While the parallel and scaling performance of CESM is well documented (Worley et al. 2011; Dennis et al. 2012; Evans et al. 2013), several important advances have recently been made by John Dennis (CISL) and collaborators to prepare CESM for the Intel Xeon Phi architectures. This effort was funded by an Intel Parallel Computing Center for Weather and Climate Simulations (IPCC-WACS) grant, and has been further supported under the NERSC Exa-scale Science Application Program (NESAP). This work has resulted in significant reductions in the cost to run both 1° and high-resolution versions of CESM. A recent estimate on the impact this investment in code optimization indicates that they reduced the cost to run CESM2 on Cheyenne by approximately 15 to 45%, depending on exact scientific configuration. Some of the more important optimizations in the CESM2 code base include: improvements to the SE dynamical core used by the atmosphere model; the barotropic solver used by the ocean model; and the boundary exchange operators used by both the ocean and sea-ice model.

A major focus of the previously mentioned NESAP project was the optimization of the SE dynamical core, also known as the Higher-Order Methods Modeling Environment (HOMME) for the Knights Landing architecture. This work significantly improved both the single core performance on Xeon Phi and Xeon in addition to improvements to the overall scalability. Preliminary description of this work is given in Dennis et al. (2017), while a more detailed description is available in Dennis et al. (2018). Figure 2 illustrates the impact this work had on the computational cost of HOMME on both Cori and Cheyenne. The computational cost in node hours per simulated year both before and after the NESAP support optimization work is provided for a range of core counts at 100 km resolution. The figure illustrates that the optimization effort reduced the cost of the HOMME dynamical core by nearly 50% on Cori, and 19 to 50% on Cheyenne depending on node count. For an ideally scaling application, the cost to run a simulation would be constant. It is interesting to note that the cost on Cheyenne tends to vary considerably with core count. It is thought that this decrease in computational cost of HOMME is a direct result of HOMME fitting into the L3 cache when more than 6 nodes are used.

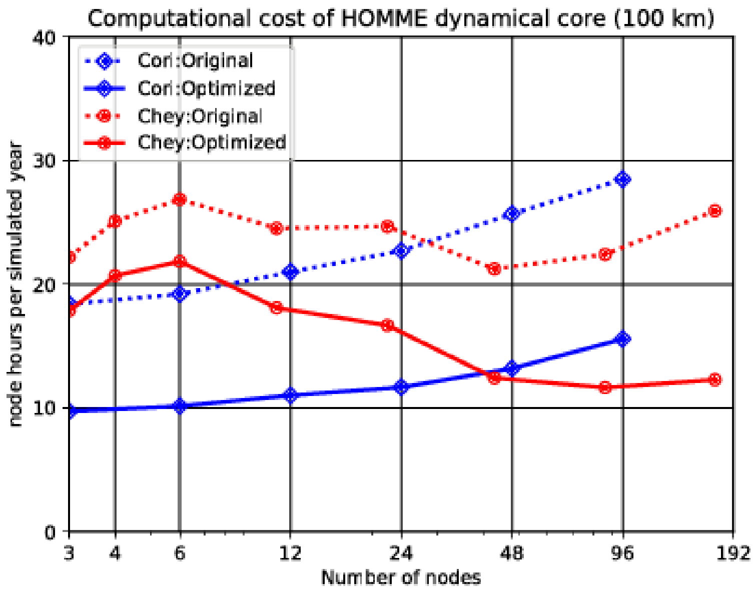

Figure 2. Computational cost of the SE dynamical core (HOMME) at 100 km resolution on both Intel Xeon and Intel Xeon Phi architectures.

Historically, the scalability of both the ocean and sea-ice model has been somewhat disappointing at larger core counts when using the 0.1° tripole grid. While the recent addition of an improved barotropic solver in the ocean model (Huang et al. 2016) has resulted in improvements, ocean model scalability issues still remained. A recent detailed investigation of the boundary exchange module used by both models revealed that a particular buffer was significantly larger than necessary. Specifically, it was $O(np^{(1/2)})$, where 'np' is the number of MPI ranks, larger than necessary. Further, every time a boundary exchange call was invoked, a very large buffer was allocated and subsequently deallocated on a small subset of the total number of MPI ranks. This over allocation created a significant load imbalance in both the ocean and sea-ice models. New boundary exchange modules, which minimized the total size of buffer, had a rather significant impact on the cost of both the sea-ice and ocean models. Specifically, the cost of the sea-ice model was reduced by about 50% at large core counts while the cost of the ocean model was reduced by about 30% on large core counts.

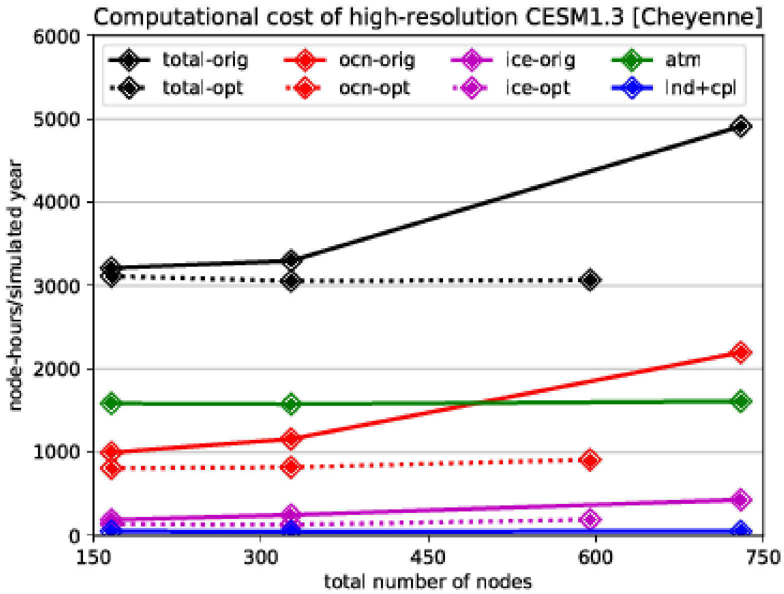

Figure 3. The computational cost of a high-resolution configuration of CESM1.3 on Cheyenne.

Figure 3 illustrates the impact that these boundary exchange modules have on the computational cost of CESM1.3 using the ne120_t12 grid on Cheyenne. Because the same improvements to the ocean and sea-ice boundary exchange module are included in the released CESM2 code, we expect a similar impact on cost at large core counts. The figure provides the total cost as well as those for the ocean, atmosphere, sea-ice, coupler, and land model. Due to the marginal cost of the coupler and land model, only their combined cost is provided for clarity. The solid lines indicate the original cost of the models while the dotted lines show the cost of the ocean and sea-ice models. Again, a perfect linear scaling would result in a completely flat cost curve. While both the optimized ocean and sea-ice models do increase in cost slightly when the node count increases from 157 to 595, the increase is very modest. Further the entire cost of all of CESM is fundamentally constant for all node counts, only varying from 3056 to 3112 node hours.

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 numerous physical, chemical, and biological processes that determine past, present, and future climates. Furthermore, they are critical 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 partnerships with scientists from universities, national laboratories, and other research organizations. CESM enables the investigations of new scientific problems through collaborations with a community continuously increasing in size, empowering many new partnerships.

A major milestone that was accomplished during the current allocation cycle was the long-waited, community release of the CESM version 2.0, CESM2.0, in early June 2018. This new version contains many substantial science and infrastructure improvements and capabilities for use of the broader CESM community. Because of the delay in the release of CESM2, the bulk of the CESM CSL priorities remain essentially the same as in the last proposal. These priorities are: i) performing CMIP6 DECK and MIPs Tier 1 simulations, primarily using the nominal 1° model version; ii) performing many of the MIPs Tier 2 and Tier 3 simulations that are particularly designed to address specific science questions that arose from the previous MIP exercises; iii) continuing development and testing of the high-resolution configurations for CESM2; and iv) continuing component model development efforts towards inclusion in the next generation model version, i.e., CESM3, as detailed in the WGs' requests, including development of the new ocean model, MOM6; re-evaluation of coupling methods between sea-ice and the ocean model; and incorporating new capabilities in CISM with the goal of supporting climate simulations with a dynamic Antarctic ice sheet. As such, this CSL request has two primary thrusts. The first is a focus on applications and pushing the frontier of resolution with the recently released CESM2 (goals # 1, 2, and 3). The second is to drive the next generation model, whose development has continued even through the period of delay in CESM2 (goal # 4).

CESM's participation in CMIP6 DECK and MIPs simulations is its broadest community project. These simulations reach a vast national and international group of scientists and researchers who rely on NCAR, and CESM in particular, to perform these simulations. Among the modeling groups which contribute to CMIP6, CESM is very unique in its community involvement and the level of transparency with which it approaches model development. CESM's contributions to the USGCRP at a national and IPCC at an international level are realized through these CMIP simulations. High level of scrutiny and analysis of the CESM CMIP6 simulations, in turn, feeds back to CESM, promoting further model development as well as enhancing collaborations.

To achieve the above priorities, the present proposal is subdivided into three sub-requests, totaling 460 M core-hours. The first is for performing all the CMIP6 DECK and MIPs Tier 1 simulations in which the CESM community will participate. These simulations are rather high priority both because they are expected to take about 10 months from now and because the manuscript submission deadline is December 2019 for inclusion in the upcoming IPCC report. The simulations require about 147 M core-hours, corresponding to about 1/3 of our two-year request. The second sub-request for order 43 M core-hours is for *Community Projects*, containing simulations that are of interest to multiple WGs, and thus, to the broader CESM community. The third sub-request is for WGs' development and production projects, totaling about 270 M core-hours over two years. These proposals request time for various model development activities towards CESM3; MIPs Tier 2 and Tier 3 simulations that are of scientific interest to the WG members; significant global and regionally-refined high-resolution development efforts as well as their evaluation; and many curiosity-driven research activities aimed at improving our scientific understanding of the Earth's climate system.

We are projecting to generate about 10.1 PB of data. In order to minimize the usage of HPSS for storing output, we request a total of 6.6 PB on campaign storage of which 4 PB will be specifically dedicated to storage of the CMIP6-related data sets. To help CISL partially cover the costs associated with this CMIP6-related storage, the CESM program requested and obtained funds from the NCAR directorate for CISL to purchase 4 PB of campaign storage. Developments in CDG that allow data distribution via campaign storage instead of HPSS will be used as much as possible. Furthermore, to minimize the usage of HPSS for storing development results, we request 1.2 PB of CESM GLADE project space. This will allow for the post-processing of community project integrations. At the end of their order 5-year life time on campaign storage, the majority of the CMIP6-related data sets will likely be archived on HPSS. However, it is anticipated that this HPSS need will be much less than the 6.3 PB produced by the CMIP6-related simulations. To the extent possible, the data from other WG and community projects simulations will be kept on campaign storage with only data sets deemed necessary archived on HPSS.

The CSL computer resources remain 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 WGs and the SSC. 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 resources 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, the tables of individual WG development and production simulations presented in the Supplementary Material include a prioritization of proposed simulations in the event that a reduced allocation is awarded.

The level of computational and storage resources requested here will allow i) timely participation of CESM2 in CMIP6; ii) performing simulations that are of scientific interest to the CESM community, also providing diagnostics on model performance in all aspects of an Earth system; iii) exploration of cutting-edge global and regionally-refined high-resolution configurations; and iv) continuation of the development path towards CESM3.

References

Bacmeister, J. T., and Co-authors, 2014: Exploratory high-resolution climate simulations using the Community Atmosphere Model (CAM). *J. Climate*, **27**, 3073-3099.

Dennis, J. M., and Co-authors, 2012: CAM-SE: A scalable spectral element dynamical core for the Community Atmosphere Model. *Inter. Jour. for High Perf. Computer Appli.*, **26**, 74-89.

Dennis, J. M., and Co-authors, 2017: Preparing the Community Earth System Model for Exascale Computing, in *Exascale Scientific Applications: Scalability and Performance Portability*, CRC publishers.

Dennis, J. M., B. Dobbins, C. Kerr, and Y. Kim, 2018: Optimizing the HOMME dynamical-core for multi-core platforms. *Inter. J. High Performance Computing Applications* (submitted).

Evans, K. J., P. H. Lauritzen, S. K. Mishra, R. B. Neale, M. A. Taylor, and J. J. Tribbia, 2013: AMIP simulation with the CAM4 spectral element dynamical core. *J. Climate,* **26**, 689-709.

Eyring, V., S. Bony, G. A. Meehl, C. A. Senior, B. Stevens, R. J. Stouffer, and K. E. Taylor, 2016: 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.

Gent, P. R., and Co-authors, 2011: The Community Climate System Model version 4. *J. Climate*, **24**, 4973-4991, doi: 10.1175/2011JCLI4083.1.

Huang, X., and Co-authors, 2016: Accelerating the high-resolution ocean model component in the Community Earth System Model. *Geosci. Mod. Dev.*, **9**, 4209-4225.

Hurrell, J., and Co-authors, 2013: The Community Earth System Model: A framework for collaborative research. *Bull. Amer. Meteor. Soc.*, **94,** 1339-1360, doi: 10.1175/BAMS-D-12-00121.1.

Kay, J. E., and Co-authors, 2015: The Community Earth System Model (CESM) Large Ensemble project: A community resource for studying climate change in the presence of internal climate variability. *Bull. Amer. Meteor. Soc.*, **96**, 1333-1349, doi: 10.1175/BAMS-D-13-00255.1.

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.

Otto-Bliesner, B. L., and Co-authors, 2016: Climate variability and change since 850 CE: An ensemble approach with the Community Earth System Model. *Bull. Amer. Meteor. Soc.*, **97**, 735-754, doi: 10.1175/BAMS-D-14-00233.1.

Small, R. J., and Co-authors, 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., and Co-authors, 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.

Worley, P. H., A. P. Craig, J. M. Dennis, A. A. Mirin, M. A. Taylor, and M. Vertenstein, 2011: Performance and Performance Engineering of the Community Earth System Model. *Proceedings of the ACM/IEEE Inter. Conf. for High Perf. Computing, Networking, Storage and Analysis (SC11)*, Seattle, WA, November 14-18, 2011.

Yeager, S. G., and Co-authors, 2018: Predicting near-term changes in the Earth System: A large ensemble of initialized decadal prediction simulations using the Community Earth System Model. *Bull. Amer. Meteor. Soc.* (in press).