


Overview and Update on the Energy Exascale Earth System Model Project


Dave Bader, Ruby Leung,
Mark Taylor, Renata McCoy
E3SM Executive Committee



Work performed under the auspices of the US
DOE by LLNL under Contract
DE-AC52-07NA27344. LLNL-PRES-xxxxxxx

EESM PI Meeting
August 2024







Overview of the E3SM Project

David C. Bader, E3SM Council Chair and
Lead Principal Investigator

August 6, 2024





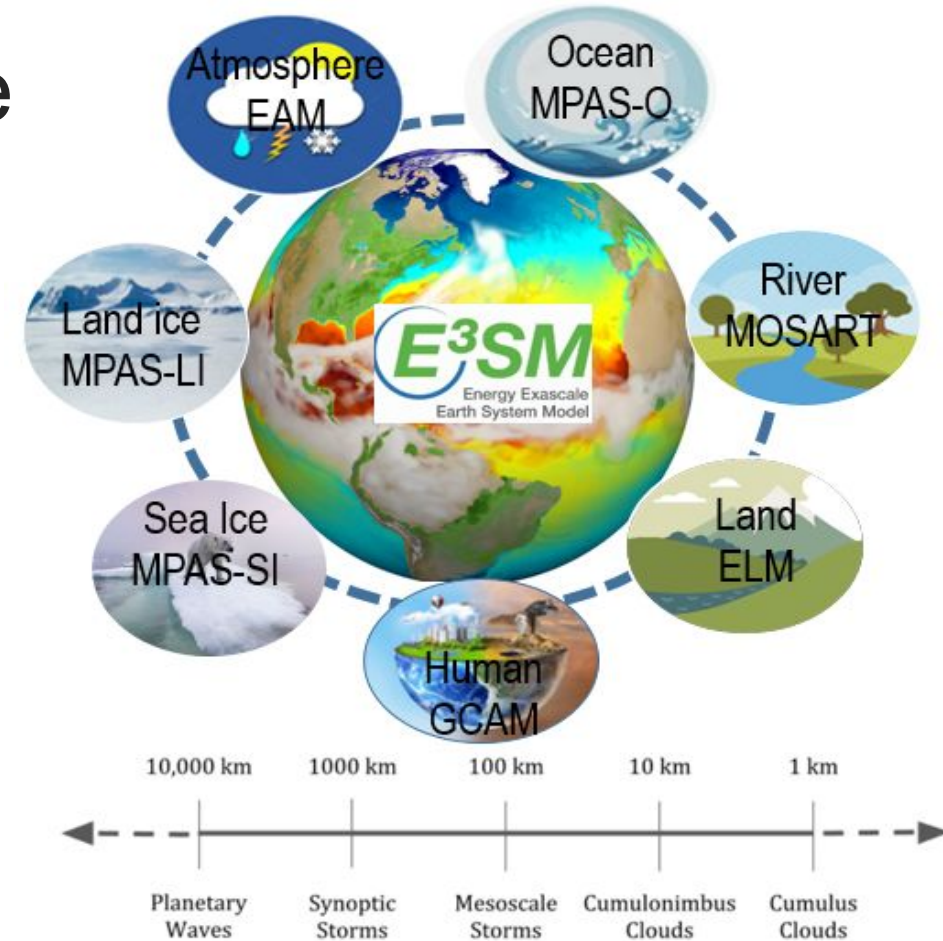
The E3SM Mission: Use exascale computing to carry out high-resolution Earth system modeling of natural, managed and man-made systems, to answer pressing problems for the DOE.*





***The E3SM project's long-term goal is to assert and maintain international scientific leadership in the development of Earth system models that address the grand challenge of actionable modeling and projections of Earth system variability and change, with an emphasis on addressing the most critical challenges facing the nation and DOE.**

The E3SM Project is at the center of an ecosystem for DOE research. The project:

- Develops the model and enables further development by others;
- Supports basic and applied DOE research requiring climate and Earth system models;
- Executes simulation campaigns and supports a modeling system to enable simulation by other DOE researchers; and
- Participates in major national and international modeling activities to ensure that the model is recognized as state-of-the-science, e.g. CMIP-DECK+.







E3SM is delivering an Exascale modeling system. Strategic Planning for the next 10 years is underway


- Pushing past past exascale will require ever-more ***disruptive approaches*** such as edge computing, machine learning (ML), and next-generation artificial intelligence (AI) to accelerate the fusion of observations and measurements with computing.
- The E3SM project will ***continuously integrate advanced technologies*** and Earth system science to deliver capabilities for multi-resolution modeling of the coupled human–Earth system.
- E3SMv4 will be at the center of a connected scientific ecosystem for understanding and modeling the Earth system, and will be ***the foundation for digital twins of the system and its components.***
- E3SM will lead in actionable projections of human–Earth system evolution across a broad range of time and spatial scales to ***support DOE’s energy mission.***

We are here to listen and learn


E3SM Science

Ruby Leung, PNNL

E3SM Chief Scientist



U.S. DEPARTMENT OF
ENERGY



Energy Exascale
Earth System Model


E3SM actionable science goals

- High-resolution modeling of extreme weather events in a changing climate
- Represent natural, managed and manmade systems and their interactions to project future outcomes
- Ensemble modeling to quantify uncertainty

Earth system science



Computational science



Modeling across scales in three versions over a decade

Beyond v3: unification

Model component	Lower resolution (LR)	High resolution (HR)	Cloud-resolving (SCREAM)	Regional refined model (RRM)
Atmosphere & Land	100 km	25 km	3 km	variable
Ocean & Ice	30-60 km	6-18 km	prescribed	variable
River	50 km	12 km	3 – 12 km	variable

**CMIP6 DECK, C4MIP
LENS, DAMIP**

HighResMIP

DYAMOND


CMIP6 DECK (NARRM)

North America RRM

25 km → 100 km




14 km → 60 km





Southern Ocean RRM

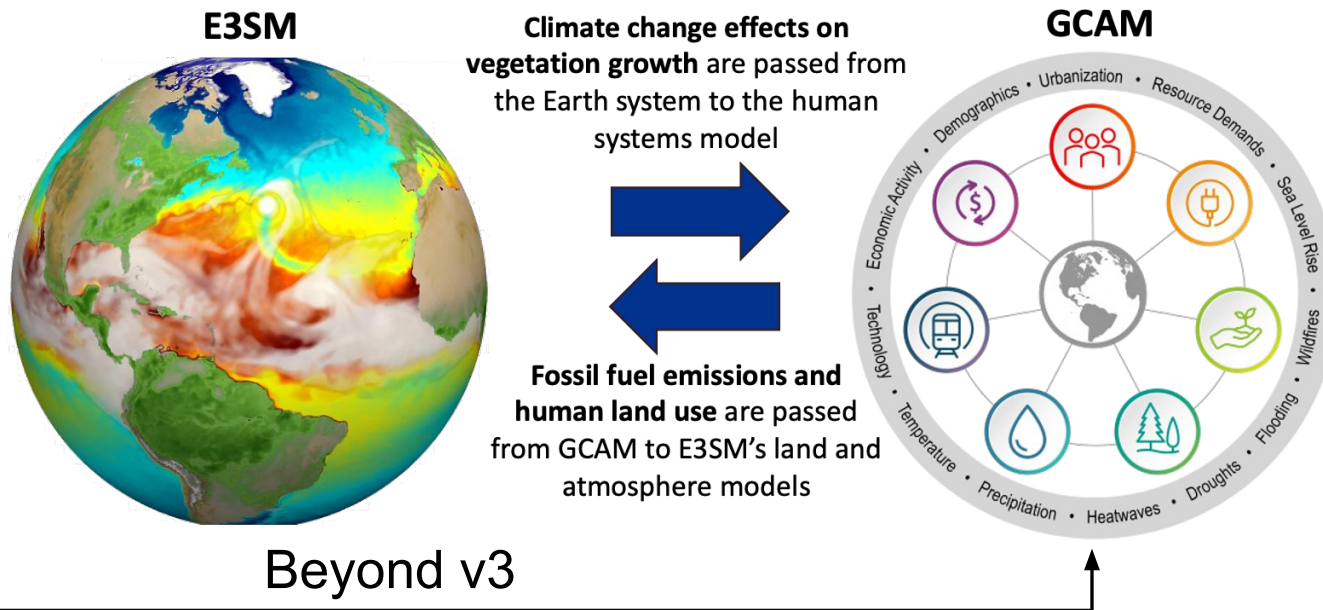
25 km → 100 km



12 km in the Antarctic, 30-60 km elsewhere



Representing human-Earth interactions



Evaluate feasibility and impacts of decarbonization scenarios



Ensemble modeling to quantify uncertainty

- In collaborations with universities and NCAR, E3SMv1 and v2 have been used to produce initial condition large ensemble simulations
- E3SMv3 is being used to produce large ensemble simulations with ~ 100 ensemble members

Three variants of
E3SMv3 with high,
medium, low **ECS**

X

25 **initial condition**
ensemble members with
a high forcing scenario

+



25 initial condition
ensemble members with
medium ECS and a low
forcing scenario

- Beyond v3, address uncertainties related to model resolution across scales, emission-driven simulations and projections, human-Earth interactions
- Other possibilities: use AI emulators for ensemble boosting to quantify uncertainty and to generate more scenarios trained using E3SM and E3SM-GCAM

E3SM Computation



Mark Taylor, SNL

E3SM Chief Computational Scientist



<https://www.top500.org/>

Rank	System	Cores	Rmax (PFlop/s)	Rpeak (PFlop/s)	Power (kW)
1	Frontier - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 20Hz, AMD Instinct MI250X, Slingshot-11, HPE DOE/SC/Oak Ridge National Laboratory United States	8,679,504	1,194.00	1,679.82	22,703
2	Aurora - HPE Cray EX - Intel Exascale Compute Xeon CPU Max 9670 52C 2.40Hz, Intel Data Center Max, Slingshot-11, Intel DOE/SC/Argonne National Laboratory United States	7,630,816	1,042.01	537.21	29,899
3	Eagle - Microsoft NDv5, Xeon Platinum 8480C 48C 2.0Hz, NVIDIA H100, NVIDIA Infiniband NDR, Microsoft Microsoft Azure United States	7,630,816	1,042.01	537.21	29,899
4	Supercomputer Fugaku - Supercomputer Fugaku, A64FX 48C 2.20Hz, Tofu interconnect D, Fujitsu RIKEN Center for Computational Science Japan	7,630,816	1,042.01	537.21	29,899
5	LUMI - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 20Hz, AMD Instinct MI250X, Slingshot-11, HPE EuroHPC/CSC Finland	2,752,704	379.00	537.21	29,899
6	Leonardo - BullSequana XH2000, Xeon Platinum 8358 32C 2.60Hz, NVIDIA A100 SXM4 64 GB, Quad-rail NVIDIA HDR100 Infiniband, EVIDEN EuroHPC/CINECA Italy	1,824,768	238.00	537.21	29,899
7	Summit - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148.00	537.21	29,899
8	MareNostrum 5 ACC - BullSequana XH3000, Xeon Platinum 8440Y 40C 2.30Hz, NVIDIA H100 64GB, Infiniband NDR200, EVIDEN EuroHPC/BSC Spain	680,960	138.20	265.57	2,560
9	Eos NVIDIA DGX SuperPOD - NVIDIA DGX H100, Xeon Platinum 8480C 56C 3.80Hz, NVIDIA H100, Infiniband NDR400, Nvidia NVIDIA Corporation United States	485,888	121.40	188.65	
10	Sierra - IBM Power System AC922, IBM POWER9 22C 3.10Hz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM / NVIDIA / Mellanox DOE/NNSA/LNL United States	1,572,480	94.64	125.71	7,438
11	Sunway TaihuLight - Sunway MPP, Sunway SW26010 260C 1.450Hz, Sunway, NRCC National Supercomputing Center in Wuxi China	10,649,400	93.01	125.44	15,371
12	Perlmutter - HPE Cray EX 235n, AMD EPYC 7763 64C 2.45GHz, NVIDIA A100 SXM4 40 GB, Slingshot-11, HPE DOE/SC/LBNL/NERSC United States	888,832	79.23	113.00	2,945




AMD Instinct™ MI250 Accelerator

- DOE SC Computing Landscape: Nearly all GPU based:
 - 1200 PF: OLCF Frontier (37K AMD GPUs)
 - 1000 PF: ALCF Aurora (64K Intel GPUs)
 - 79 PF: NERSC Perlmutter (6K NVIDIA GPUs, 3K CPU nodes)
- Key mission for E3SM is to run on these architectures
- Programming Model:
 - Difficult & time consuming to port E3SM to run efficiently on GPUs
- SCREAM approach:
 - Rewrite code from scratch in C++/Kokkos
 - Strong emphasis on testing
- C++/Kokkos programming model
 - Robust support across all GPUs to date
 - on-node parallelism to be expressed in vendor's native GPU language
 - Does not rely on lagging vendor support for Fortran

SCREAM



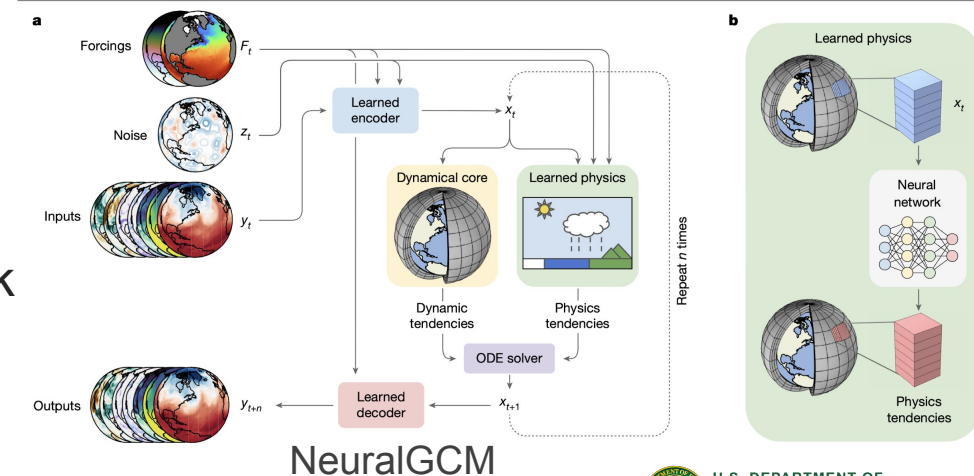

- SCREAM (evolving into EAMxx for E3SM V3)
 - Demonstrated true performance portability:
 - Competitive performance on CPUs compared to Fortran code
 - Excellent results on NVIDIA GPUs (V100s, A100s) and AMD GPUs (MI250s)
- First-to-Exascale opening up new science:
 - Broke the long standing “1 SYPD” goal for a global cloud resolving model
 - Multi-decadal length simulations at cloud resolving resolutions to be completed in 2024!
 - **2023 Gordon Bell Prize in Climate Modelling for innovative parallel computing contributions toward solving the global climate crises.**
- OMEGA: Apply the SCREAM model for E3SM V4 Ocean:
 - Clean-sheet-of-paper rewrite
 - C++/Kokkos for performance portability on CPU and GPU



Taylor, Caldwell, Bertagna, Clevenger, Donahue, Foucar, Guba, Hillman, Keen, Krishna, Norman, Sreepathi, Terai, White, Wu, Salinger, McCoy, Leung, Bader, *The Simple Cloud-Resolving E3SM Atmosphere Model Running on the Frontier Exascale System* SC23: International Conference for High Performance Computing, Networking, Storage and Analysis (2023)

Future Plans



- Run on new and upcoming GPUs:
 - Aurora's Intel GPUs: Demonstrate continued performance portability
 - Upcoming integrated CPU/GPU/HBM from NVIDIA (GH) and AMD (MI300)
- What's the future of GPUs and CPUs?
 - ML/AI now drives semiconductor development and dwarfs HPC
 - Example: 2024 AI demand has caused NVIDIA H100 nodes ~15x more expensive than CPU nodes, but prices are dropping)
 - Ultra low precision for deep neural networks
- ML and AI:
 - Phenomenal success in NWP, trained from reanalysis
 - Climate modeling: Harder due to uncertainties baked into future climate data
 - New research using ML to improve climate models and make them more efficient



E3SM Resources

Renata McCoy, LLNL

E3SM Chief Operating Officer & Project Engineer



- E3SM is an **Open Development Code!!**
 - As of the release of the E3SM version 1.0 (Apr 2018) E3SM's developmental code is freely available on GitHub
 - <https://github.com/E3SM-Project>
- All **E3SM tools** are also **open-source** development tools
 - <https://docs.e3sm.org/>
 - <https://e3sm.org/resources/tools/>

- **Data from all major simulations is published on Earth System Grid Federation (ESGF)**

- Released E3SM Data



- **Only “Scientifically Validated” configurations are supported**
 - maint-1.0 , maint-1.1, maint-1.2, maint-2.0, maint-2.1,
 - Guaranteed to run “out-of-the-box” on the **E3SM-supported DOE LC centers**

Early Access to data, simulations or internal documentation

Submit a **collaboration request**



- Please **respect E3SM right to first publication on new developments** and wait with your publication until E3SM has a chance to submit a paper on its new development

See <https://e3sm.org/resources/policies/>




Documentation

- Centralized E3SM Model Documentation
 - <https://docs.e3sm.org/>
- Extensive documentation on public confluence
 - <https://acme-climate.atlassian.net/wiki/spaces/DOC>
- Online Tutorials
 - <https://e3sm.org/publications/events/e3sm-tutorials/>




2024 E3SM Tutorial



E3SM Communication

- E3SM Website
 - <http://e3sm.org>
- E3SM Conferences and Presentations
 - <https://acme-climate.atlassian.net/wiki/spaces/ECM>
- E3SM quarterly "Floating Points" Newsletter provides:
 - Latest news
 - Research Highlights
 - Project vision and Roadmaps
 - Self subscribe:
 - email_listserv@listserv.llnl.gov with the email body: 'subscribe E3SM-news'
- E3SM YouTube Channel: seminars, webinars, and tutorials
 - <https://www.youtube.com/@e3sm-project>

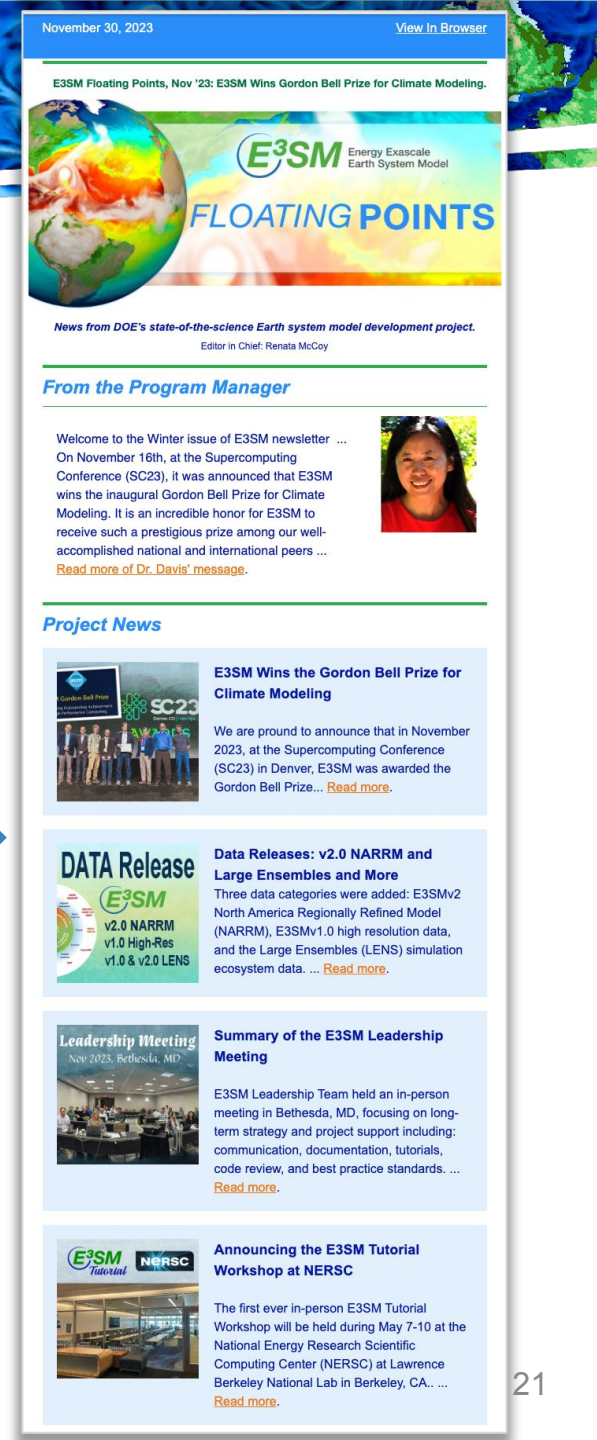


Newsletter Archive



For more Info

– Joint **Networking Session**: "E3SM Tutorial/Informational Session", **today at 4pm**



e3sm.org

