#### **AGU Town Hall**

E3SM: A Decade of Earth System Modeling Effort at the Department of Energy

Xujing Davis Asmeret Asefaw Berhe Gary Geernaert Dave Bader Ruby Leung Mark Taylor Renata McCoy

#### E3SM Wins Gordon Bell Prize for Climate Modeling, SC23









ESMD: Earth System Model Development Program Area



# DOE Office of Science Director Perspective: Role of E3SM for DOE Mission and Science

Asmeret Asefaw Berhe DOE Office of Science Director







# Earth System Modeling Priorities at DOE and Interagency Landscape

Gary Geernaert DOE EESSD Director



EESSD: Earth and Environmental Systems Sciences Division





# Program Overview Xujing Jia Davis DOE ESMD Program Manager



ESMD: Earth System Modeling Development Program Area





# **ESMD** in Support of E3SM



#### Innovative and computationally advanced ESM capabilities, in support of Energy science and mission

Earth System Across Scales



**Goal:** Support the development of E3SM including its subcomponents, to address the grand challenge of actionable predictions of the changing Earth system, emphasizing on the most critical scientific questions facing the nation and DOE

#### Strategies:

- Science driver for model development
- Earth system across scales (high-resolution frontier, bridge gaps, quantify uncertainty via LE)
- Prepare for and overcome the disruptive transition to next era of computing, leverage ASCR HPC capabilities
- Innovative mathematical, computational methods, tools, algorithms (e.g., ML/AI)

**EAM:** E3SM Atmosphere Model; **ELM:** E3SM Land Model ; **GCAM:** Global Change Assessment Model; **MOSART:** Model for Scale Adaptive River Transport; **MPAS-SI:** Model for Prediction Across Scales (MPAS) – Sea Ice; **MPAS-O:** MPAS – Ocean; **MALI:** MPAS-Albany Land Ice Model.

More E3SM Acronyms: https://e3sm.org/resources/help/acronyms/



#### ESMD Portfolio & Ecosystem in Support of E3SM

**E3SM:** An integrator of DOE earth, environmental, mathematical and computational sciences, in advancing ESM capability for DOE science mission.

FY 23 Budget Distribution



#: U. FOA, CEDS, Interagency ... ...

#### Note: Univ. scientists across ESMD projects

See detail about ESMD Projects

FY 23: \$49 M

#### **ESMD** supported Projects:

- > Funding instruments:
- 1. Lab-led projects including Scientific Focus Area (SFAs, e.g., E3SM); 2. Scientific Discovery through Advanced Computing (SciDAC) Awards; 3. Early Career Awards and 4. Other projects: e.g., U. FOA, Interagency activities (e.g, USGCRP/IGIM, <u>CICE</u> <u>Consortium</u>...)
- E3SM SFA is the central driver of the E3SM development with focused scientific questions, well defined time frames, goals and strategies
- > Other projects contribute to E3SM development in various ways on different time frames

#### E3SM in DOE ecosystem EESSD:

- *RGMA*: PCMDI, RUBISCO, HYPERFACETS, WACCEM, HILAT-RASM, CATALYST, CASCADE...
- ➢ MSD: GCIMS (GCAM), HYPERFACETS, IM3 …
- > ARM/ASR: Field Campaigns, THREADS, LASSO ....
- > ESS: NGEE-Arctic, NGEE-Tropics, SPRUCE, COMPASS-FME, Urban IFL ...

ASCR: SciDAC, Exascale Computing Project (ECP) ...

Office of Science: Energy Earthshot, RENEW, FAIR, RDPP, CRC ...





#### Science Community Leadership and Service

#### National

- > USGCRP: IGIM US Climate Modeling Summit (USCMS), GEWEX: D. Bader, R. Leung
- > NASEM Digital Twin Workshop: R. Leung, M. Taylor
- > NCA5: R. Leung, P. Thornton, C. Tebaldi, P. Ullrich
- ➤ US CLIVAR : R. Leung
- OSTP ICAMS Subcommittee on Earth System Modeling and Prediction(ESM&P) Implementation Teams: *M. Taylor, R. Jacob, C. Golaz, P. Jones, A. Donahue, O. Guba*
- > CESM Advisory Committee: *E. Hunke, M. Taylor*

#### International

- CICE Consortium: E. Hunke, A. Roberts
- International CLIVAR: L. Van Roekel
- > International Workshop on Coupling Technologies for ESMs: *R. Jacob*
- > WCRP GEWEX Global Atmospheric System Studies Panel (GASS) annual meeting: S. Xie
- Association for Computing Machinery (ACM) and the Swiss National Supercomputing Centre: O. Guba



E3SM contributes to national and global endeavor in advancing Earth System Predictability while addressing the DOE mission





023 E3SWI AU-Hands Meeting, Denver, CO

#### **E3SM Timeline and Major Achievements**



**Coupled Earth-Human Feedback**: coupling with GCAM





#### E3SM Leadership Team: Cross Laboratory Initiative

#### Speakers today→



■ David Bader, Chair ■ Ruby Leung, Chief Scientist ■ Mark Taylor, Chief Computational Scientist ■ Renata McCoy, Project Engineer



















E3SM Council Chair and Lead Principal Investigator









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.





# E3SM Approach



- Major simulations. A series of simulation-andprojection experiments addressing mission needs with actionable scientific results.
- **Model development.** A well-documented, tested, continuously improving system of model codes that comprise the E3SM Earth system model.
- Leadership architectures. The ability to use effectively leading (and "bleeding") edge computational facilities soon after their deployment at DOE national laboratories.
- Infrastructure. An infrastructure to support code development, hypothesis testing, simulation execution, and analysis of results.







# E3SM Phase 2 Highlights

- E3SM is possible because of a strong culture of "project before lab," and the commitment of talented and dedicated scientists, computational scientists and software engineers.
- Completion of v1 Simulation Campaign
- Over 12,000 simulated years of simulations using v2 E3SM
- v2 RRM with consistent model tuning with v2 standard resolution
- Atmosphere algorithmic improvements doubled model throughput
- Demonstrated templated C++ programming model for hybrid CPU/GPU (Exascale) machines that requires little support from compiler vendors
- Established a more rigorous Code Review/Testing process to enable more predictable integration of new developments, both internal and external, eg EAGLES.
- Installed and maintained E3SM system on NERSC and COMPY computers for use by other BER/EESSD programs





The Simple Cloud-Resolving E3SM Atmosphere SCREAM GCRM (3.25 km) Benchmark Performance Full Model Model (SCREAM)

- DOE has the fastest computers in the world, but they use NVIDIA, AMD, and Intel GPUs
  - Weather/climate models require major modification to run on GPUs
  - No single programming strategy works for all 3 GPU vendors
- ⇒ "Performance portability" was needed for E3SM to achieve its exascale ambitions
  - This was achieved by writing a new atmosphere model in C++/Kokkos
  - SCREAM won the 2023 Gordon Bell Climate Prize for breaking 1 simulated year per day at Δx=3.25 km



Fig: throughput vs node count at  $\Delta x=3.25$ km on Frontier (AMD GPUs) and Summit (NVIDIA GPUs)





#### The Phase 3 Concept



Conceptual diagram of parallel development paths





# E3SM is on the verge of delivering an Exascale modeling system. What's next?

- 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*.
- DOE will lead in actionable projections of human–Earth system evolution across a broad range of time and spatial scales to *support multisectoral decision making and DOE's energy mission*.



### E3SM Science Ruby Leung, PNNL E3SM Chief Scientist









#### Overarching goal: advance actionable science in support of DOE's energy mission

Climate change impacts on energy supply, delivery, and demand



(NCA5 Chapter 5)

Climate Change Impacts on the Energy System

#### **Science drivers:**

- Water cycle changes and impacts
- Human-Earth system feedbacks
- Polar processes, sea-level rise, and coastal impacts







# 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







# Modeling across scales

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, C4MI	P HighResMIP	DYAMOND	CMIP6 DECK (NARR	



Southern Ocean RRM

Delaware Bay RRM

4 km → 240 km

25 km → 100 km

















### **Regional refinement: NARRM**

NARRM has similar climate sensitivity as LR, simplifying model calibration

![](_page_22_Figure_4.jpeg)

Improved orographic precipitation

#### **Reduced SWCF bias from stratocumulus**

![](_page_22_Figure_6.jpeg)

#### More realistic tropical cyclone tracks

![](_page_22_Figure_8.jpeg)

Simulated TC tracks

![](_page_22_Figure_10.jpeg)

![](_page_22_Picture_11.jpeg)

(Tang et al. 2023 GMD)

![](_page_23_Picture_0.jpeg)

#### A new unified surface mesh over land, river, and ocean

![](_page_23_Figure_2.jpeg)

![](_page_23_Picture_3.jpeg)

Use of tri-grid and unstructured mesh in all model components provides flexibility for telescoping to the grid spacing needed to model multiple flood drivers and their interactions in coastal regions

![](_page_23_Picture_5.jpeg)

![](_page_24_Picture_0.jpeg)

![](_page_24_Picture_1.jpeg)

### Modeling human-Earth system feedbacks

![](_page_24_Figure_3.jpeg)

![](_page_25_Picture_0.jpeg)

#### Modeling human-Earth system feedbacks

![](_page_25_Figure_2.jpeg)

![](_page_25_Picture_3.jpeg)

![](_page_26_Picture_0.jpeg)

# Modeling polar processes: ice shelf melt fluxes

#### Most climate models

![](_page_26_Figure_3.jpeg)

![](_page_26_Figure_4.jpeg)

![](_page_26_Figure_5.jpeg)

Southern Ocean Regionally Refined Mesh (SORRM) resolution (km)

![](_page_26_Figure_7.jpeg)

#### **Configuration:**

- Resolution of 12 km in the Antarctic, ~30-60 km elsewhere
- Prognostic ice shelf melt fluxes, data iceberg melt climatology

C21D-1262 (Asay-Davis et al.) on Tuesday, 9:30am – 1:30pm Melt fluxes in control, two ensemble members of historical and future (SSP370) simulations

![](_page_26_Figure_13.jpeg)

![](_page_27_Picture_0.jpeg)

# Modeling polar processes: Arctic sea ice

![](_page_27_Figure_2.jpeg)

ENERGY 28

![](_page_28_Picture_0.jpeg)

# E3SM Exascale Readiness Mark Taylor, SNL

E3SM Chief Computational Scientist mataylo@sandia.gov

![](_page_28_Picture_3.jpeg)

![](_page_28_Picture_4.jpeg)

![](_page_28_Picture_5.jpeg)

EXASCALE COMPUTING PROJECT

![](_page_29_Picture_0.jpeg)

# **DOE Computing Landscape**

- DOE: Large investments in Exascale computing
  - E3SM mission to run on these computers
  - DOE SC machines ranked #1, #2, #7, #12 in the Top500 list of the world's fastest computers
  - Most of the compute power, power consumption and purchase cost comes from the GPUs
- E3SM Exascale mission = run efficiently on GPUs
- Challenge: GPUs can provide large acceleration of many kernels, but what about the full model, including all the time spent in communication?
  - E3SM's global cloud resolving atmosphere (SCREAM)
  - 5.8x faster (AMD 2x64 core CPU node vs. 4xA100 GPU node)
  - 3.5x faster per Watt (based on measured power consumption)

#### 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 2GHz, AMD Instinct MI250X, Slingshot-11, HPE DOE/SC/Oak Ridge National Laboratory United States	8,699,904	1,194.00	1,679.82	22,703
2	Aurora - HPE Cray EX - Intel Exascale Compute Blade, Xeon CPU Max 9470 52C 2.4GHz, Intel Data Center GPU Max, Slingshot-11, Intel DOE/SC/Argonne National Laboratory United States	4,742,808	585.34	1,059.33	24,687
3	Eagle - Microsoft NDv5, Xeon Platinum 8480C 48C 26Hz, NVIDIA H100, NVIDIA Infiniband NDR, <b>Microsoft</b> Microsoft Azure United States	1,123,200	561.20	846.84	
4	Supercomputer Fugaku - Supercomputer Fugaku, A64FX 48C 2.20Hz, Tofu interconnect D, Fujitsu RIKEN Center for Computational Science Japan	7,630,848	442.01	537.21	29,899
5	LUMI - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE EuroHPC/CSC Finland	2,752,704	379.70	531.51	7,107
6	Leonardo - BullSequana XH2000, Xeon Platinum 8358 32C 2.6GHz, NVIDIA A100 SXM4 64 GB, Quad-rail NVIDIA HDR100 Infiniband, EVIDEN EuroHPC/CINECA Italy	1,824,768	238.70	304.47	7,404
7	Summit - IBM Power System AC922, IBM POWER9 22C 3.076Hz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148.60	200.79	10,096
8	MareNostrum 5 ACC - BullSequana XH3000, Xeon Platinum 8460Y+ 40C 2.3GHz, 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 3.8GHz, NVIDIA H100, Infiniband NDR400, Nvidia NVIDIA Corporation United States	485,888	121.40	188.65	
10	Sierra - IBM Power System AC922, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM / NVIDIA / Mellanox DOE/NNSA/LLNL United States	1,572,480	94.64	125.71	7,438
11	Sunway TaihuLight - Sunway MPP, Sunway SW26010 260C 1.456Hz, Sunway, NRCPC National Supercomputing Center in Wuxi China	10,649,600	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	888,832	79.23	113.00	<sup>2,945</sup> 30

![](_page_30_Picture_0.jpeg)

# Programming Models for GPU systems

#### Fortran + Directives

- Relies heavily on (lagging) vendor compiler support
- Remains immature w.r.t. advanced Fortran features
- · Good performance still requires major code refactoring

#### C++ / on-node Parallel array abstractions ( Kokkos and YAKL)

- C++ has robust vendor support across NVIDIA, AMD and Intel
- Kokkos and YAKL backends quickly adapt to each vendor's preferred technology (CUDA, HIP, SYCL, pthreads, etc...)
- Kokkos and YAKL rely on standard C++ methodology and work together seamlessly

#### DSL: Domain Specific Language

- Promising approach being explored by several modeling centers (e.g. GT4Py, GridTools, PSyclone)
- Most HPC experience within DOE labs is with C++

![](_page_30_Picture_13.jpeg)

![](_page_30_Picture_14.jpeg)

![](_page_30_Picture_15.jpeg)

![](_page_30_Picture_16.jpeg)

![](_page_31_Picture_0.jpeg)

# E3SM GPU Capabilities

- SCREAM: Simple Cloud Resolving E3SM Atmosphere Model
  - New nonhydrostatic atmosphere model, rewritten from scratch in C++
  - Designed for cloud resolving resolutions (prescribed aerosols, no deep convection parameterizations)
  - Competitive with Fortran version on CPUs, and running well on NVIDIA, AMD (and hopefully soon Intel) GPUs
  - 2023 Gordon Bell Prize for Climate Modeling!

#### • E3SM-MMF (superparameterized-E3SM)

- E3SM fully coupled simulations at typical climate resolution, running with most atmosphere parameterizations replaced by a local cloud resolving model
- GPU acceleration allows MMF approach to obtain similar throughput as a conventional model
- In Progress:
  - Omega: Port of MPAS-Ocean into C++/YAKL
  - ELM: GPU acceleration via Fortran/OpenACC (transitioning to OpenMP to support AMD and Intel GPUs)

![](_page_31_Picture_13.jpeg)

![](_page_32_Picture_0.jpeg)

# GPU enabled simulations Roadmap

- Cloud resolving atmosphere simulations
  - 2023: Multi-year simulations for Cess-Potter climate sensitivity
  - 2024: Multi-decadal AMIP simulations
- E3SM-MMF
  - 2023: Century long full coupled climate simulations with cloud resolving deep convection "super-parameterization"
- E3SMv4: 2026
  - Full Earth System Model running efficiently on GPUs
  - Atmosphere: based on SCREAM, with additional support for non-cloud resolving resolutions, prognostic aerosols and some chemistry
  - Ocean: Omega eddy resolving ocean
  - Less expensive components (land, lce): use idle CPU cores of the GPU node, or run on GPUs depending on readiness

![](_page_32_Figure_12.jpeg)

![](_page_32_Picture_13.jpeg)

![](_page_32_Picture_14.jpeg)

# E3SM as an Open Science Model Development Project Renata McCoy, LLNL

E3SM Chief Operating Office & Project Engineer

![](_page_33_Picture_2.jpeg)

![](_page_33_Picture_3.jpeg)

![](_page_34_Picture_0.jpeg)

![](_page_34_Picture_1.jpeg)

# **Open Science Development Model**

- E3SM is **Open Development Code**!!
  - As of the first data release (Apr 2018), E3SM is freely available on GitHub
- All in E3SM tools are open-source development tools
  - https://e3sm.org/resources/tools/
- Data from all major, campaign simulations are available to all, published on ESGF \*
  - Few months delay between the production and publication to ESGF to publish an overview paper
- Support for code and data (limited)
  - Support for "scientifically validated" compsets / configuration used in simulation campaigns
  - Supported versions:
    - <u>maint-v1.0</u>, <u>maint-v1.1</u>, <u>maint-v1.2</u>, <u>maint-v2.0</u>, <u>maint-v2.1</u>,
  - Guaranteed to run on the E3SM-supported DOE LC centers

![](_page_34_Picture_14.jpeg)

![](_page_35_Picture_0.jpeg)

![](_page_35_Picture_1.jpeg)

### External project's early access to data or internal documentation

Projects can submit a **collaboration request** specifying:

- What data/simulation/early access information they need
- What research are they planning to do
- Specify an E3SM Point of Contact (POC)
- Agree to collaborate and include the E3SM POC in your publication
- Fill in the doc at
  - <u>https://e3sm.org/about/collaboration/collaboration\_request/</u>

![](_page_35_Picture_10.jpeg)

![](_page_36_Picture_0.jpeg)

# Documentation, Webinars, Online Tutorials

- Extensive documentation on the Confluence public site
  - <u>https://acme-climate.atlassian.net/wiki/spaces/DOC</u>
- Diagnostics and analysis tools
  - <u>https://e3sm.org/resources/tools/</u>
- Documentation and online tutorials
  - <u>https://e3sm.org/about/events/e3sm-tutorials/</u>
- Webinars and presentations
  - <u>https://e3sm.org/about/events/all-hands-presentations/</u>

![](_page_36_Picture_11.jpeg)

![](_page_36_Picture_12.jpeg)

![](_page_37_Picture_0.jpeg)

# First In-person Hands-on Tutorial Workshop @ NERSC

- May 7-10, 2024
- At NERSC
  - National Energy Research Scientific Computing Center (NERSC) Lawrence Berkeley National Lab in Berkeley, CA.
- Applications
  - Announce via E3SM and EESSD-related email lists in December
  - Priority to DOE projects
  - <u>https://e3sm.org/announcing-the-e3sm-tutorial-workshop-at-nersc</u>
- The tutorial will encompass:
  - 1. Lectures on earth system simulation and the model components of E3SM.
  - 2. Practical sessions on running E3SM, modifying components, and analyzing data.
  - 3. Best practices for utilizing the model and potentially contributing to its development.
- Will be recorded and available on E3SM YouTube
  - Lectures, examples, class notes

![](_page_37_Picture_15.jpeg)

![](_page_37_Picture_16.jpeg)

![](_page_37_Picture_17.jpeg)

![](_page_38_Picture_0.jpeg)

# E3SM Communication

- E3SM Website <a href="http://e3sm.org">http://e3sm.org</a>
- Public Confluence for
  - Developer's documentation: https://acme-climate.atlassian.net/wiki/spaces/DOC
  - E3SM Conferences: 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 with seminars, webinars, and tutorials
  - https://www.youtube.com/@e3sm-project

![](_page_38_Picture_14.jpeg)

![](_page_38_Picture_15.jpeg)

![](_page_38_Picture_16.jpeg)

![](_page_38_Picture_17.jpeg)

Welcome to the Winter issue of E3SM newslette On November 16th, at the Supercomputing Conference (SC23), it was announced that E3SM wins the inaugural Gordon Bell Prize for Climate Modeling. It is an incredible honor for E3SM to receive such a prestigious prize among our well accomplished national and international peers Read more of Dr. Davis' message

![](_page_38_Picture_19.jpeg)

#### **Project News**

![](_page_38_Picture_21.jpeg)

#### E3SM Wins the Gordon Bell Prize for Climate Modeling

We are pround to announce that in Novembe 2023, at the Supercomputing Conference (SC23) in Denver, E3SM was awarded the ordon Bell Prize Read mor

![](_page_38_Picture_24.jpeg)

#### Data Releases: v2.0 NARRM and Large Ensembles and More Three data categories were added: E3SMv2 North America Regionally Refined Model (NARRM), E3SMv1.0 high resolution data, and the Large Ensembles (LENS) simulation

![](_page_38_Picture_26.jpeg)

#### Summary of the E3SM Leadership

E3SM Leadership Team held an in-persor eeting in Bethesda, MD, focusing on long erm strategy and project support including communication, documentation, tutorials code review, and best practice standard Read more

![](_page_38_Picture_29.jpeg)

#### Announcing the E3SM Tutorial Vorkshop at NERSC

The first ever in-person E3SM Tutorial Workshop will be held during May 7-10 at the National Energy Research Scientific Computing Center (NERSC) at Lawrence erkeley National Lab in Berkeley, CA.. Read more

39

![](_page_39_Picture_0.jpeg)

![](_page_39_Picture_1.jpeg)

E3SM-related sessions on AGU: https://tinyurl.com/a8yvy958