

EarthWorks Progress

David Randall
AMWG meeting February 2025





EarthWorks



EarthWorks is a five-year project, in which CSU and 3 NCAR laboratories are working together as partners.

The goal is to develop a *global convection-permitting and ocean-eddy permitting coupled model* based on the CESM.

Earthworks is supported by a grant from NSF/CISE. The CSU/NCAR split is 60/40.

EarthWorks is also partnering informally with DOE, through our use of MPAS-Ocean and interactions with Andrew Gettelman and scientists at the Los Alamos National Lab.



This is a team effort.




David Randall¹, James Hurrell¹, Donald Dazlich¹, Lantao Sun¹, Andrew Feder¹,
William Skamarock², Brian Medieros², Xingying Huang²,
Sheri Mickelson², Supreeth Suresh², Thomas Hauser², Ming Chen²,
Dylan Dickerson², Brian Dobbins², Michael Duda², Jim Edwards², Chris Fisher²,
Jihyeon Jang², Mariana Vertenstein², John Clyne², Richard Loft³, Andrew Gettelman⁴,
Phil Jones⁵, Luke Van Roeckel⁵, John Cazes⁶, Gunther Huebler⁷,
Raghu Raj Prasanna Kumar⁸, Pranay Reddy Kommera⁸, Carl Ponder⁸

¹Colorado State University, ²National Center for Atmospheric Research, ³AreandDee, LLC, ⁴Pacific Northwest National Laboratory,
⁵Los Alamos National Laboratory, ⁶University of Texas at Austin, ⁷University of Wisconsin, Milwaukee, ⁸NVIDIA

NCAR software engineers have been key to our progress.

Rich Loft and Sheri Voelz created some of the slides in this talk.




Special thanks to Don Dazlich, Brian Dobbins, Jim Edwards,
Rich Loft, and Sheri Voelz.



The EarthWorks Vision



- Using CESM as a base, create a coupled model that uses the same 3.75-km geodesic grid for atmosphere, ocean, and land surface. Use of the same grid for all components is central to the EarthWorks vision.
- Develop improved parameterizations that work well with both 3.75-km and 120-km grids.
- Use the model to study both weather and climate, on time scales ranging from days to years.

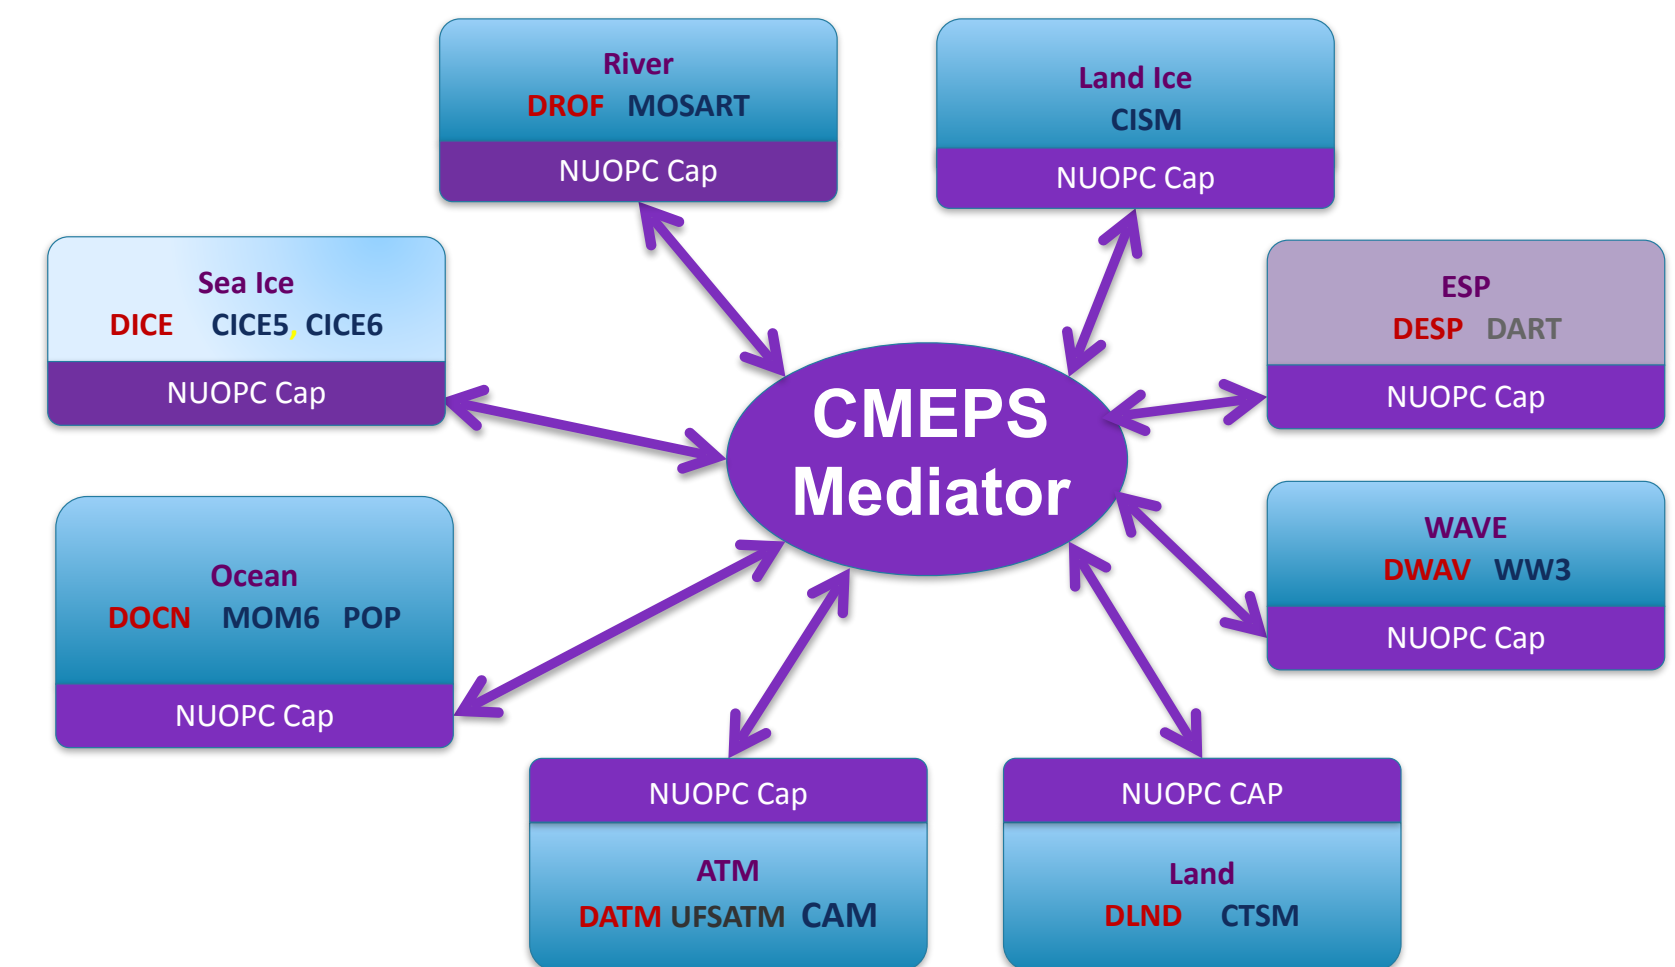




EarthWorks consists of:



- The MPAS non-hydrostatic dynamical core, with a resolved stratosphere and CAM-ish physics
- The MPAS ocean model, developed at Los Alamos
- The MPAS sea ice model, based on CICE
- The Community Land Model (CLM)
- The Community Mediator for Earth Prediction Systems (CMEPS)
- The Common Community Physics Package (CCPP), when ready





Scientific Goals




- Replace the deep convection and gravity-wave drag parameterizations with direct simulation.
- Demonstrate the resolution-dependent advantages and disadvantages of using the same grid for all ESM components.
- Use the model to study the interactions of mesoscale weather systems with larger spatial scales, on time scales ranging from days to years.
- Use lower-resolution versions of the model to study century-scale climate change.
- Perform deterministic forecasts (days to S2S) with a higher-resolution version of the same model that is used for climate-change experiments.
- Use the model to understand deficiencies of lower-resolution versions of CESM.
- Use the model to create training data sets for machine learning.




Computational Goals



- Find and fix issues with the CESM infrastructure that prevent it from being used with very high resolution.
- Provide CPU compatibility for low resolution experiments and for short tests with ultra-high resolution.
- Provide GPU compatibility (atmosphere and ocean) for ultra-high resolution simulations.
- Achieve about one SYPD with 3.75 km global grid spacing on a DOE “leadership-class” machine in 2025.






EarthWorks Version 2.3 Software Release



Release Date: Aug 1, 2024

- **GitHub Location:** github.com/EarthWorks.org/EarthWorks
- New features:
 - The RRTMG-P OpenACC data transfers were improved thanks to work by Supreeth and John Dennis.
 - Testing revealed correctness issues which were addressed by temporarily reverting to a non-OpenACC MPAS-A version....
 - Nvhpc 24.9 fixed longstanding restart issue.




EarthWorks Version 2.4 Software Release



Release Date: Dec 20, 2024

- **GitHub Location:** github.com/EarthWorks.org/EarthWorks
- New features:
 - EarthWorks is now based on the `cesm3_0_beta03` tag to onboard recent science improvements, code revisions, and changes in input data.
 - We installed a new non-hydrostatic version of MPAS-O, which has been ported to GPUs. We now have non hydrostatic dynamics for both the atmosphere and the ocean.
 - CLUBB has been ported to GPUs.
 - Other GPU improvements include support for OpenACC MPAS-A with simple physics compsets.

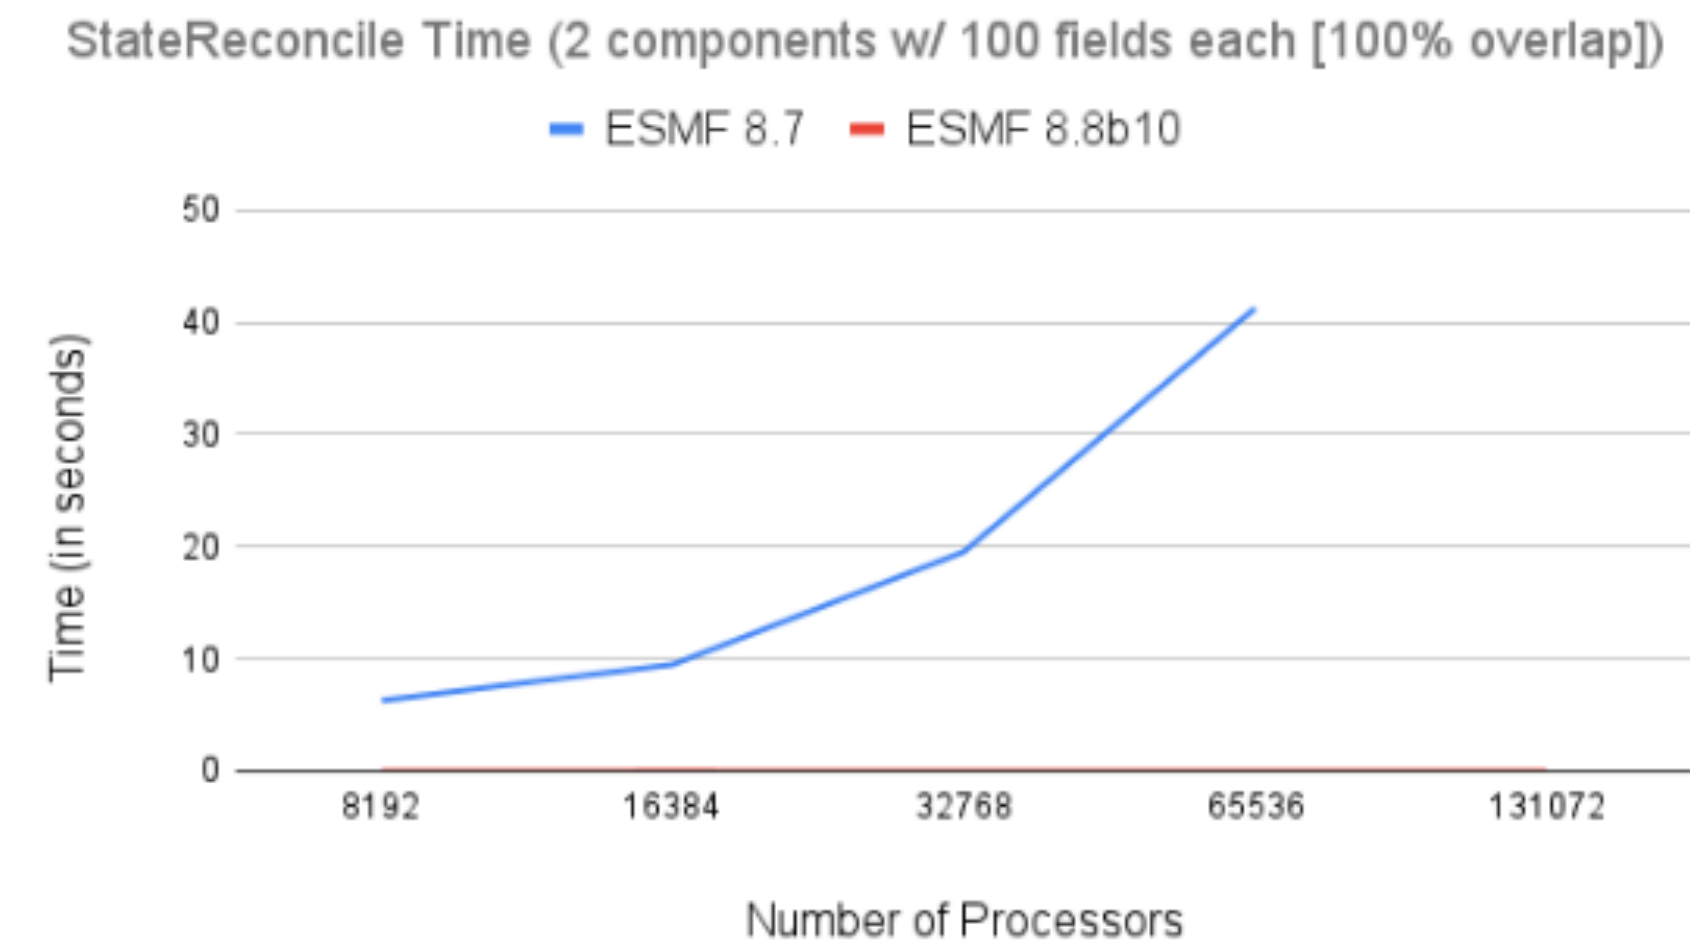


CESM Infrastructure Advances: ESMF

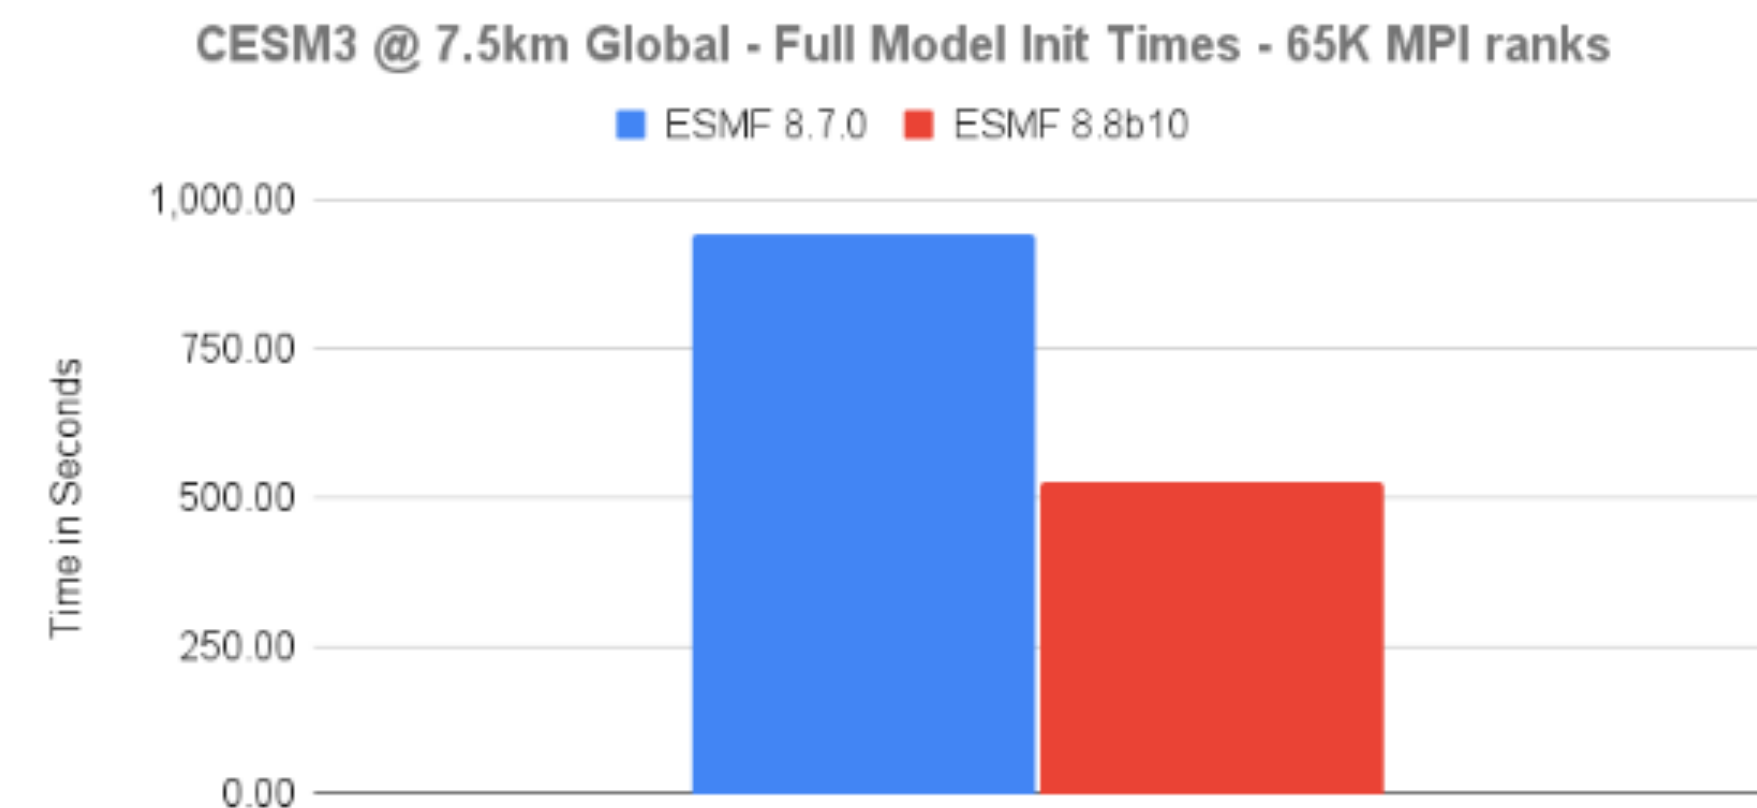



Thanks to the ESMF team

- (blue) ESMF 8.7 had scalability and performance issues with StateReconcile, which is called at initialization.
- The problem occurred even when all ES components are decomposed identically and coincident on each processing element. In fact, StateReconcile cannot even run successfully on 131K cores.
- (red) The new version of StateReconcile (ESMF 8.8 beta 10) recognizes this configuration, and uses a 'bypass approach' to achieve near-zero time for StateReconcile.



- The end-to-end initialization time is sped up by nearly a factor of 2 in a 7.5 km Atmosphere + Land test with CESM3 between ESMF 8.7 (blue) and 8,8beta 10 (red)
- In this test, all components on 65K cores. The times measured in this test include I/O, regridding, ESMF init, and component model initialization.

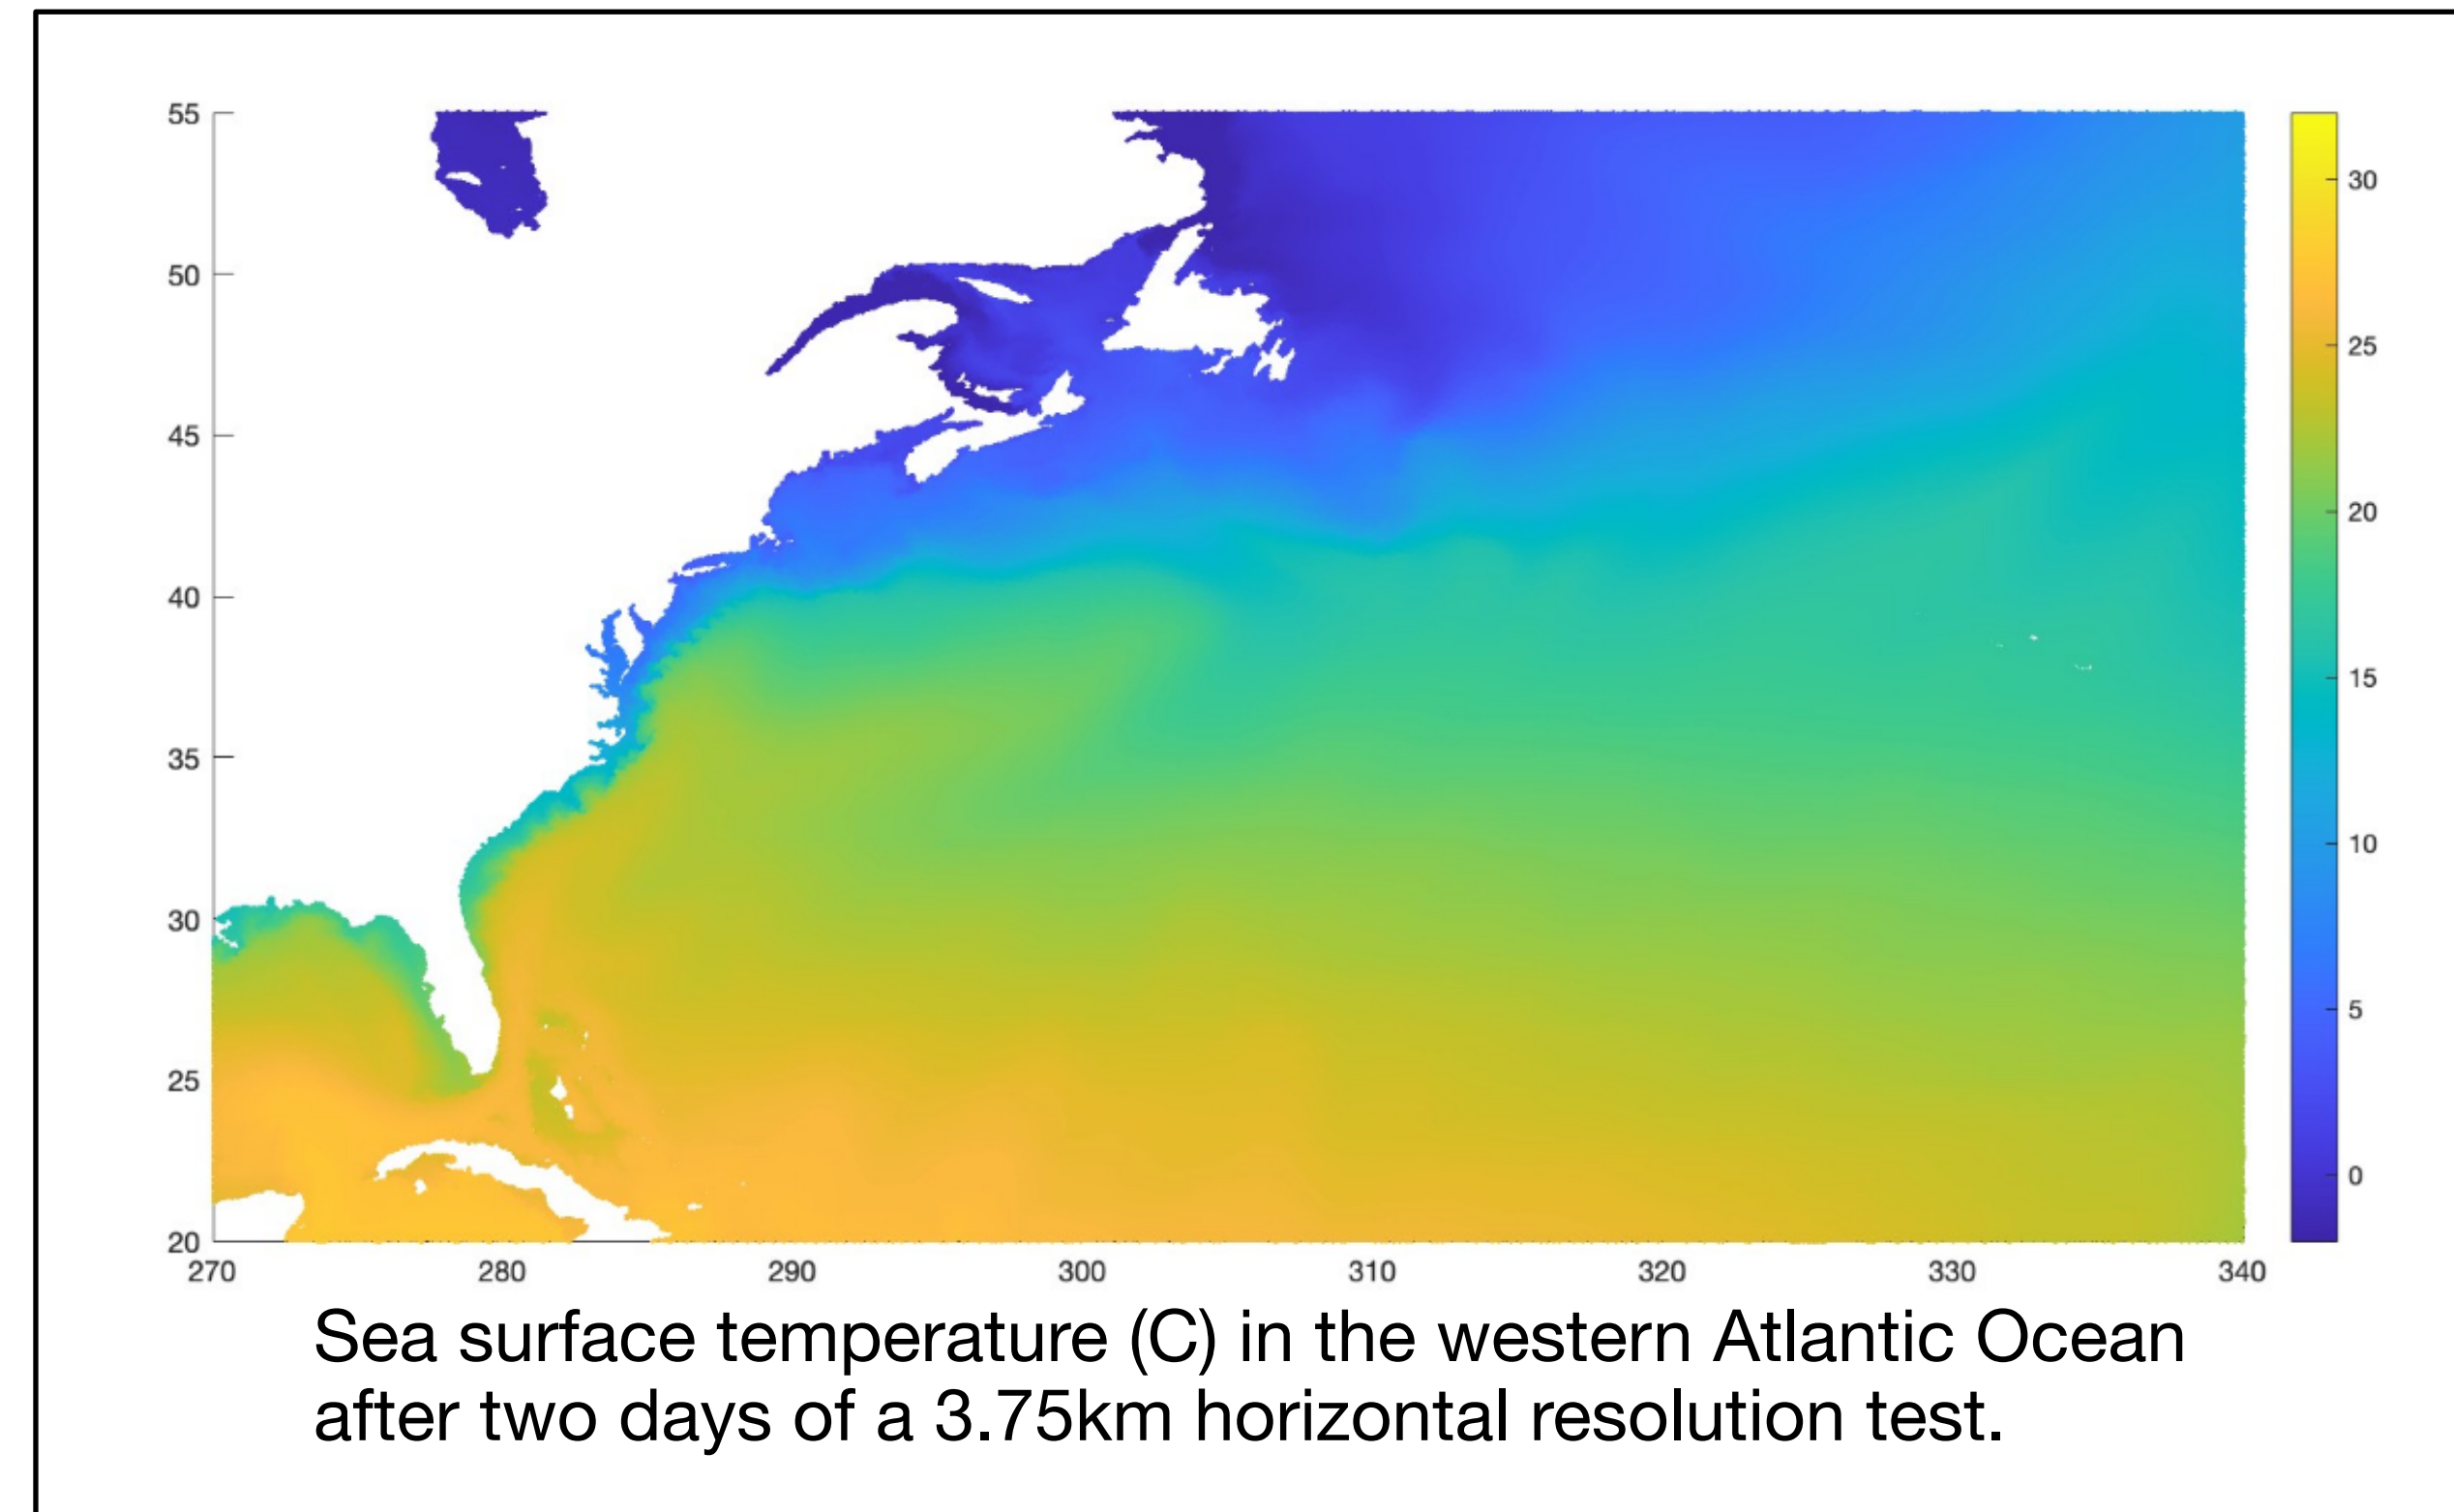




MPAS-Ocean Progress



- Updated the ocean dynamics to include a non-hydrostatic treatment of the vertical motion.
- Generated balanced and stable ocean initial conditions (T, S, layer thickness) for finer grids - 15 km, 7.5 km, 3.75 km
- Fixed bit-for-bit reproducibility on restart.
- Currently working to fix issues reading and writing ocean restarts with the 3.75-km grid.
- Next: Generate river runoff remappings for higher resolutions



CLUBB CPU/GPU Performance-Portability

Work by Gunther Huebler, grad student at the University of Wisconsin




- **Portability Comparison**

- Across offload directive paradigms
 - OpenACC vs OpenMP (via Intel conversion tool)


- **Performance Comparison**

- Across compilers (NVHPC vs Cray)
- Across Architectures
 - GPUS (Nvidia A100 vs AMD MI250X)
 - CPUs AMD “Milan” 7763 node (2 processors, 128 cores)
- Across modes: synchronous and asynchronous


Derecho - CPU / GPU Runtime Ratio (large column numbers)



Results from Gunther Huebler's presentation entitled: CLUBB GPUization and Performance Portability delivered at the 105th AMS Conference.




CLUBB CPU/GPU Performance Results – Standalone Model




**CPU optimal at
8 runs / core**

**GPU throughput
beats CPU peak
at ~25k runs**



From Gunther Huebler's presentation entitled: CLUBB GPUization and Performance Portability delivered at the 105th AMS Conference.



DYAMOND-3 'Lite'



A one-year fully-coupled simulation with a 15-km horizontal grid spacing

We have just started a simulation that follows the proposed DYAMOND-3 protocol* but with coarser resolution.

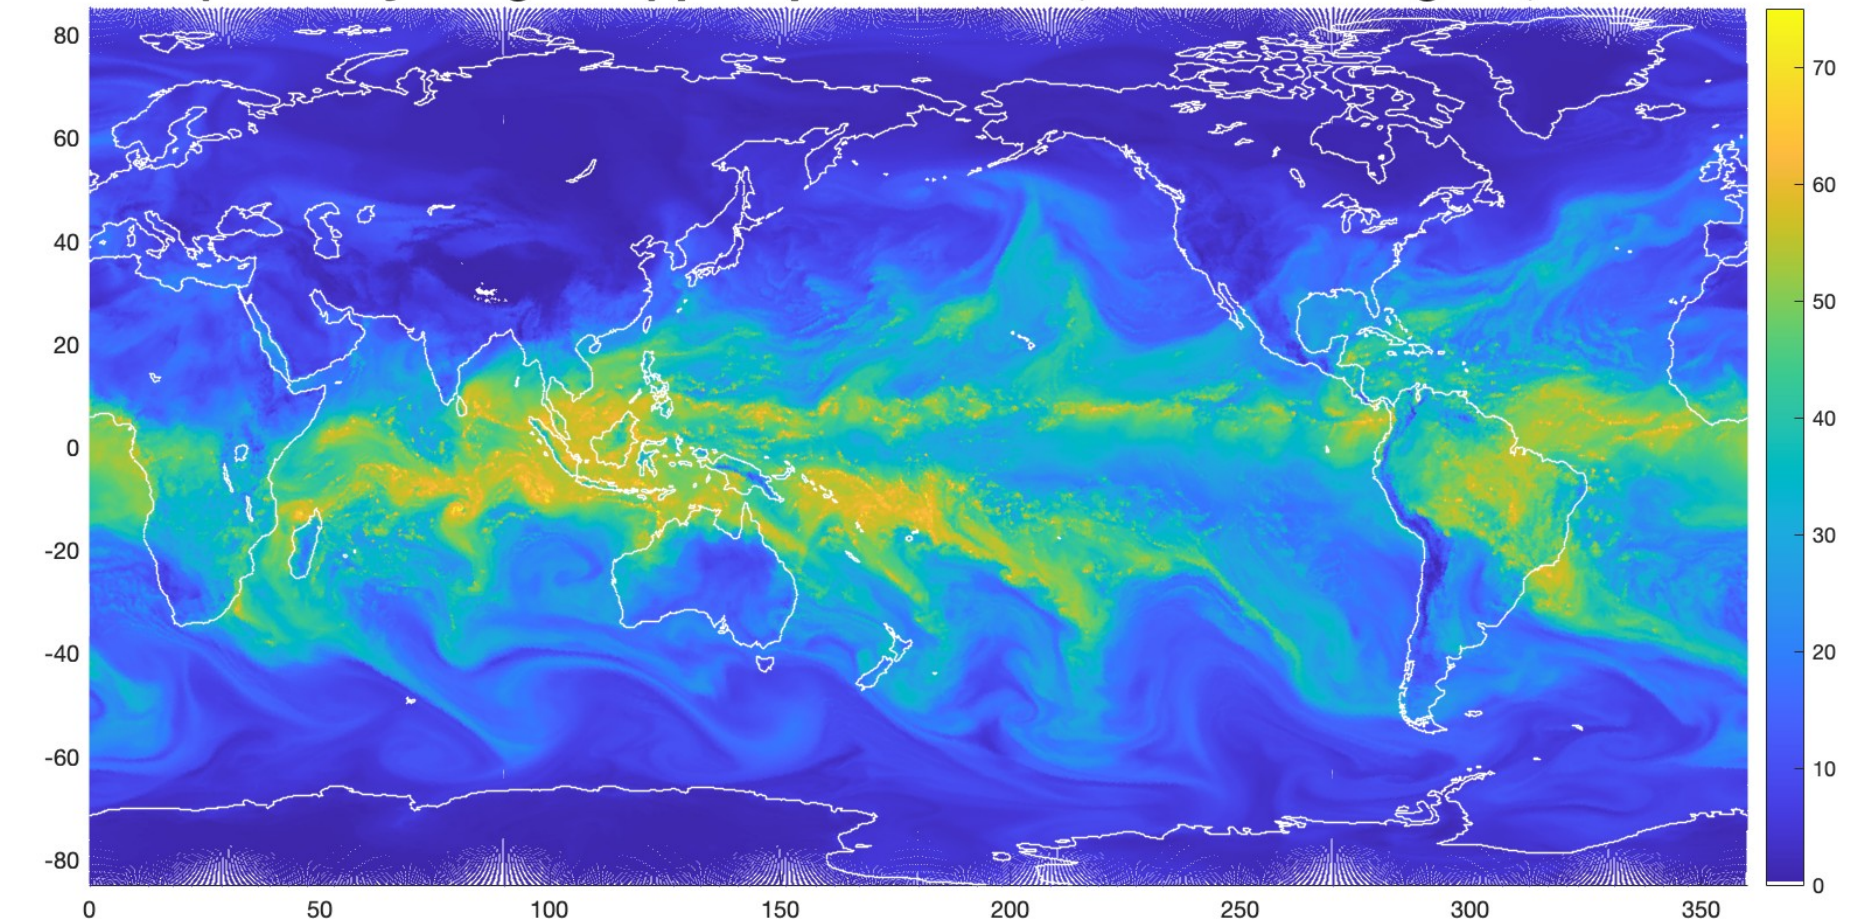
Configuration: EarthWorks v2.4, CAM 7 physics, 58 levels.

This will be the first “long” EarthWorks simulation with a grid spacing finer than 30 km.

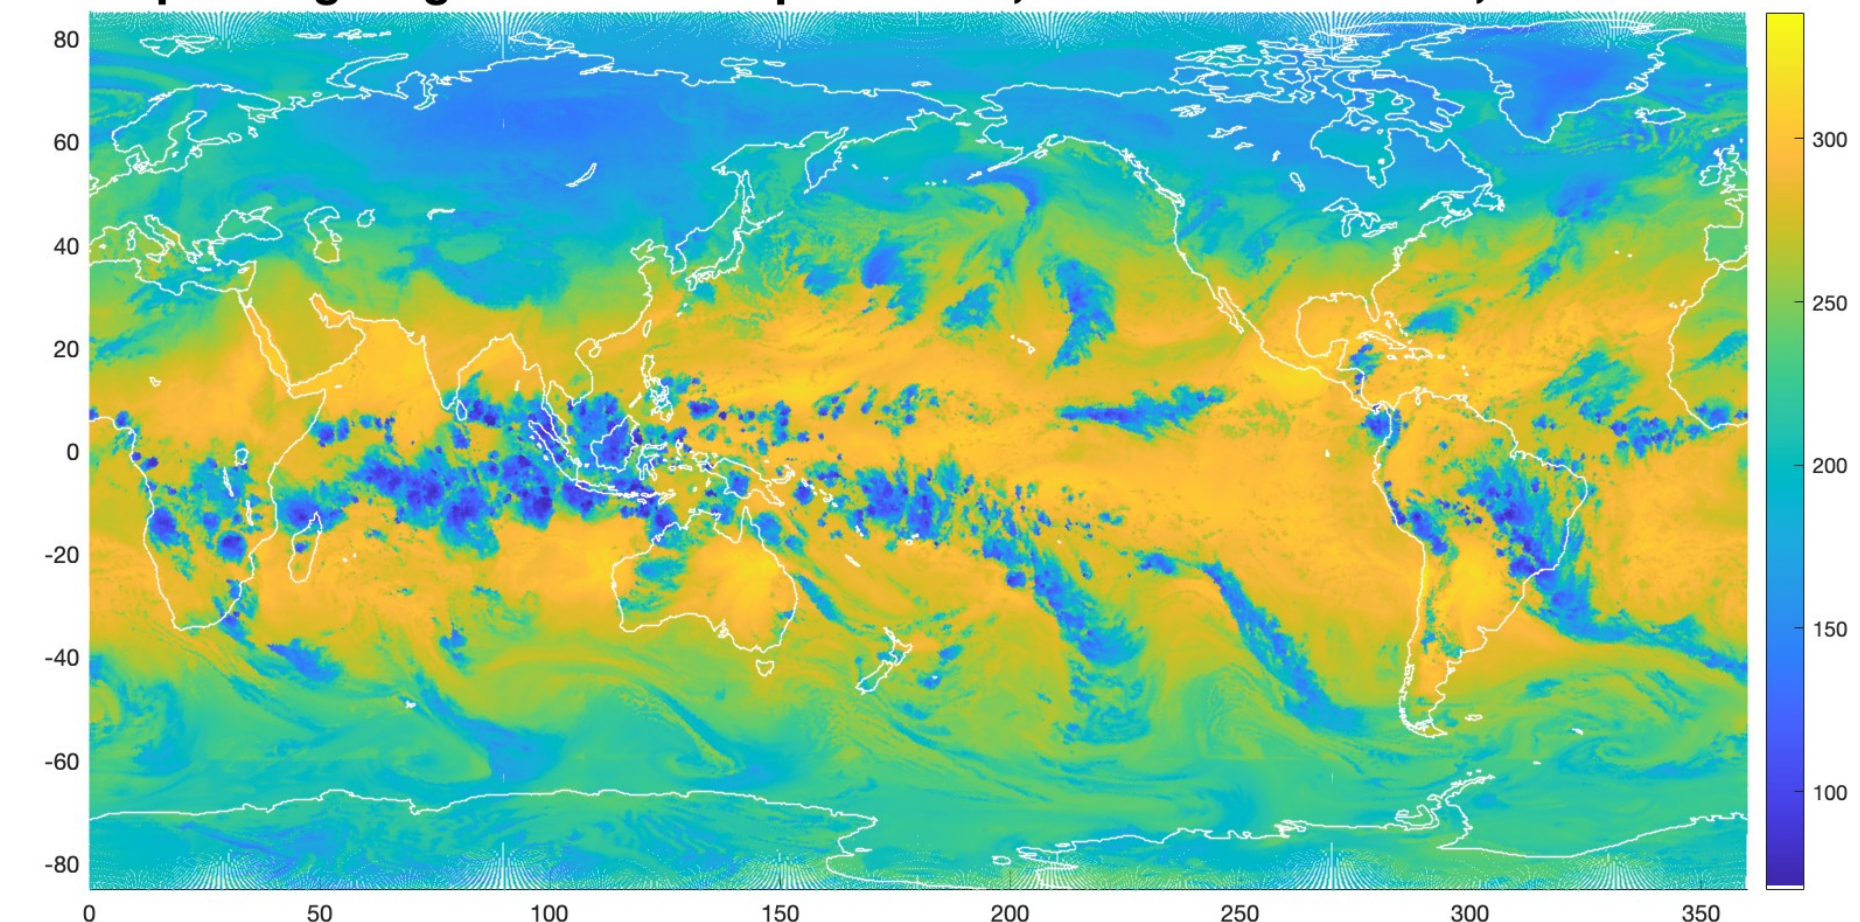
Purpose - to provide model output to develop analysis tools to process MPAS-grid generated data at the WCRP global km-scale hackathon.


*Takasuka, D., Satoh, M., Miyakawa, T. *et al.* A protocol and analysis of year-long simulations of global storm-resolving models and beyond. *Prog Earth Planet Sci* **11**, 66 (2024). <https://doi.org/10.1186/s40645-024-00668-1>

Total (vertically integrated) precipitable water, Mean =22.67kg/m2, 00Z Jan 2



Upwelling longwave flux at top of model, Mean =241.41W/m2, 00Z Jan 2







DYAMOND-I



40 days and 40 nights with a 3.75-km horizontal grid spacing (uncoupled)

We are about to start a DYAMOND I simulation.
It may already have started.






CSSI Proposal



- **Title:** *Collaborative Research Frameworks: Enhancing the speed and ease of use of a km-scale Community Earth System Model*
- **Focused on Software Engineering and Computer Science**
 - Infrastructure scalability/performance improvements
 - GPU porting (aerosols)
 - Exploring 32-bit components for faster computation.
 - Machine Learning
 - Data compression
 - Feature identification

Submitted November 22, 2024




AGS Proposal



- **Title:** *Studies of mesoscale convective systems with a km-scale Community Earth System Model*
- **Focused on Earth Science**
 - Case studies of mesoscale convective systems in Subtropical South America and the North American Great Plains
 - Interactions of mesoscale convective systems with the land surface
 - Quantifying the advantages of the common grid
 - Possibly more...

To be submitted ASAP




Illinois Collaboration



EarthWorks is collaborating with Robert “Jeff” Trapp, Kelvin Droggemeier and colleagues who have submitted a pre-proposal for an NSF Science and Technology Center.

The title of their proposed project is “Extreme Events in a Changing Climate.”

EarthWorks contributed results from a short simulation for use in the pre-proposal.



Summary



We are encountering, identifying, and fixing issues with the CESM infrastructure.

We are porting to GPUs and otherwise updating model components.

We are benefitting from collaborations.

We are performing simulations.

We are writing proposals so that the work can continue.

Thanks

