March 5 (and 6), 2026

Welcome to the 2026 CESM SEWG Winter Meeting and Hackathon

Ligia Bernardet and Michael Levy
SEWG co-chairs

Code of Conduct

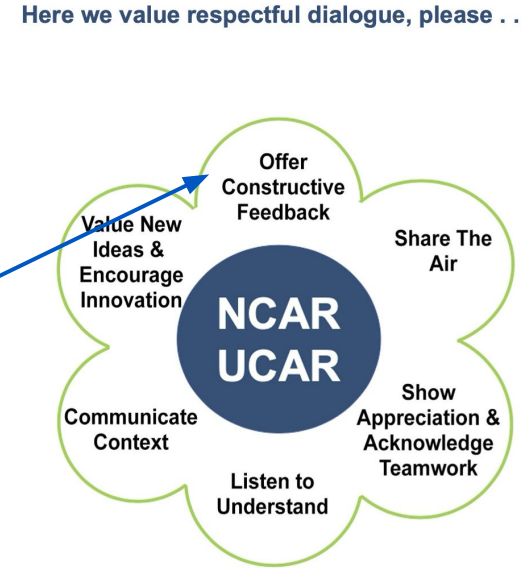

Offer Constructive Feedback: it's okay to disagree with one another, but be polite

Share the Air: make sure everyone is has a chance to participate (especially important during the hackathon!)

Agenda (big picture)

Today
9:00 - 9:30 General welcome / hackathon project pitches
9:30 - 11:45 Discussion on AI Assistants [this room] or hack time [library]
10:30 - 10:45 Break
11:45 - 1:15 Lunch
1:15 - 4:00 Discussion on running Fortran / C on GPUs [this room] or hack time [library]
2:30 - 3:00 Break
4:15 - 6:00 Happy Hour: Southern Sun (self-pay / cash-only / ATM on site)

Tomorrow *(construction outside library; possibly snow in the forecast?)*
8:30 - 11:45 Hack time: groups can spread between this room & library
11:45 - 1:15 Lunch
1:15 - 3:00 Hack time
3:00 - 4:00 Project Reports

Agenda (details)

AI Assistants

9:30 - 9:45 Sam Rabin (NSF NCAR / CGD)
9:45 - 10:00 David Appelhans (NVIDIA)
10:00 - 10:15 John Dennis (NSF NCAR / CISL)
10:15 - 10:30 Michael Kavulich (NSF NCAR / RAL)
10:30 - 11:00 BREAK
11:00 - 11:15 Aya Lahlou (Columbia; remote)
11:15 - 11:20 Kate Thayer-Calder (NSF NCAR / CGD)
11:20 - 11:45 Discussion

Agenda (details)

GPU Talks

1:15 - 1:30 Weiqun Zhang (LBL; pre-recorded)
1:30 - 1:50 Daniel Abdi (CU/CIRES at NOAA GSL; remote)
1:50 - 2:10 Ren Stengel (NSF NCAR / CISL)
2:10 - 2:30 Pranay Kommera (NVIDIA)
2:30 - 3:00 BREAK
3:00 - 3:20 Marshall Ward (GFDL)
3:20 - 3:40 John Dennis (NSF NCAR / CISL)
3:40 - 4:00 Discussion

Upcoming CESM Events

CESM Workshop: June 15 - 17, 2026

CESM Tutorial: July 6 - 10, 2026

CESM 3.0 Updates

- Finalizing configuration for PI Control
- Goal: start CMIP7 runs before June workshop
- Release date is uncertain, probably Summer 2026 (but after workshop)

With NCAR in the news a lot over the past few months, we have been told to proceed with business as usual as much as possible so that's what we are doing

June, 2026

CESM Workshop Cross Working Group Session:

Storytelling to craft climate science tools for communities

Laura Landrum, Mari Tye, Diamond Tachera, and James Done

Storytelling and Story Sharing

The importance of a good story

Stories are fundamentally how humans share information and learn.

Traditional scientific discourse is a very narrow form of storytelling.

Session aims to broaden “scientific storylines” * to include integrations of knowledge and experience

* Shepherd et al (2018). Storylines: An alternative approach to representing uncertainty in physical aspects of climate change. Climatic Change, 151(3–4), 555–571. https://doi.org/10.1007/s10584-018-2317-9

June, 2026

CESM Workshop Cross Working Group Session:

Storytelling to craft climate science tools for communities

Laura Landrum, Mari Tye, Diamond Tachera and James Done

Updates from NOAA

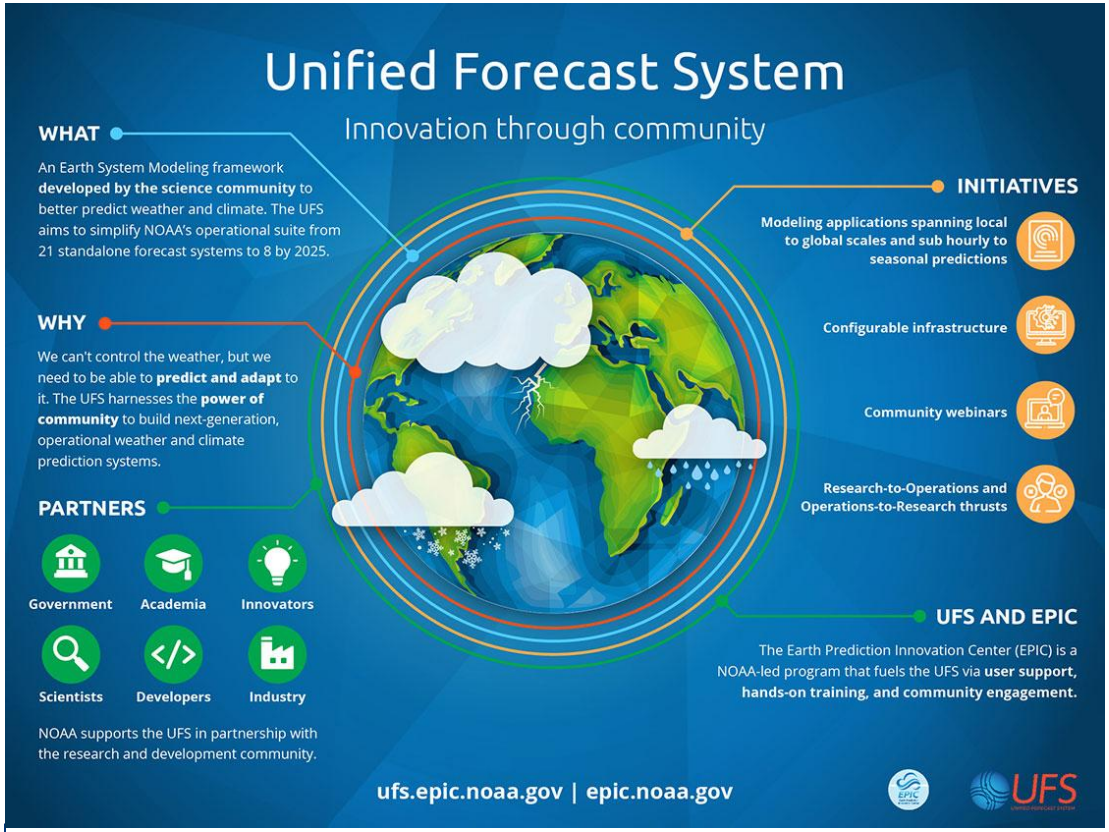

Shared between UFS and CESM

Science goals: Earth system prediction across scales

Infrastructure and (sub) Components

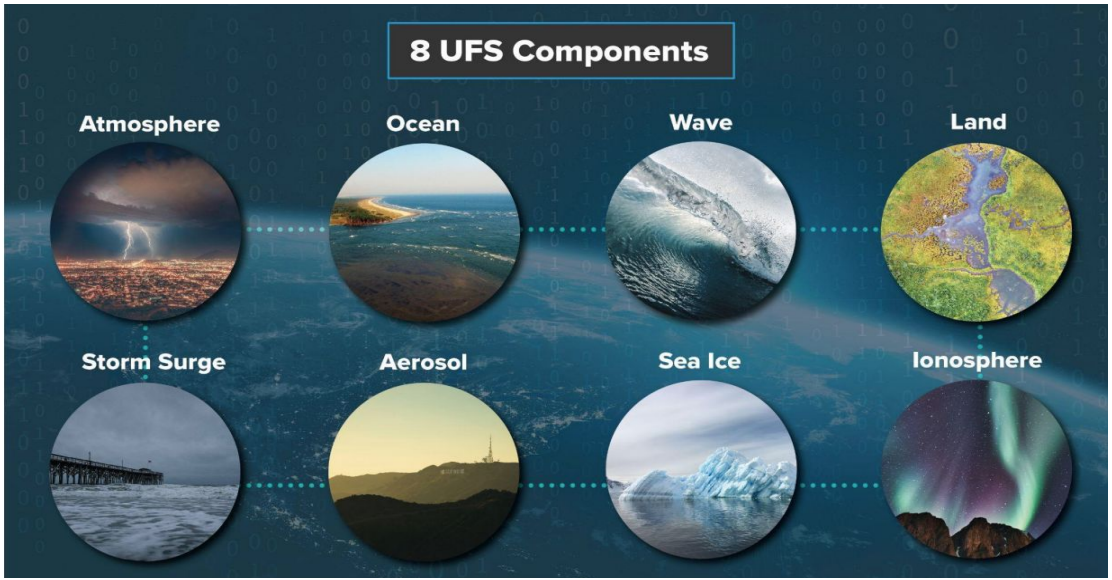

- **Atm**
  - Dycores: FV3, MPAS
  - Physics: CCPP
- **Ocean:** MOM6
- **Sea Ice:** CICE6
- **Wave:** WAVEWATCH3
- **Land:** CLM lake model
- **Chemistry:** Musica
- **Coupling**
  ESMF/CMEPS/CDEPS
- **Etc.**

Adapted from M. Huang, 2026 AMS

Recent and Upcoming Highlights

- **Hurricane Analysis and Forecast System (HAFS) v2.1** transitioned to operations in July 2025
- **Rapid Refresh Forecast System (RRFS) v1** going operational in Q4FY26
  - UFS Short Range App v3 released May 2025
  - Air Quality (with EPA CMAQ) capability added
- **Gas Chemistry and Aerosol Forecast System (GCAFS) v1** going operational in Q4FY26 (with NASA's GOCART aerosols)
- **Global Forecast System (GFS) v17** code near frozen for operational pre-test
  - First fully coupled version of GFS
- **Seasonal Forecast System (SFS) v1** under development
- And… Production suite moving to cloud computing

New: NOAA Moonshot

GOAL 1: A New Global Prediction System

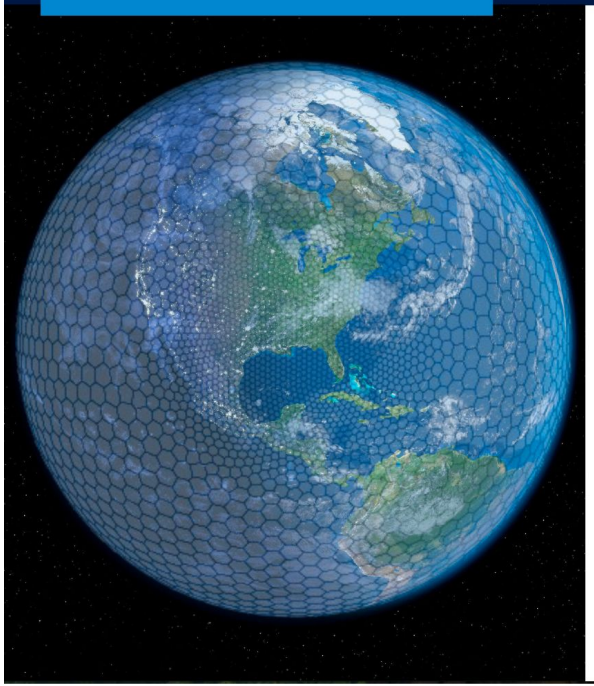

Goal 1. A global, fully-coupled and convection-allowing 3km Earth system prediction system based on the MPAS dynamical core.

5

From Bandy and Kleist, 2026 AMS

Hackathon Projects

5 hackathon projects available to work on, details to follow:

1. General CUPiD development (Google Meet: ngh-gmod-onw)
2. Running CUPiD for multi-instance CESM (Google Meet: gsz-bnwu-miw)
3. Unit Testing in Fortran / Python (Google Meet: qpx-ejub-nfe)
4. Using AMReX to Modernize MOM6 (Google Meet: atu-capf-kdo)
5. Using an AI Assistant (Google Meet: hks-bpge-wdp)

Google meet links will also be pinned in the main chat; attendance today may be light given the two sessions of talks.

Project organizers:

- *Make sure meet is open so virtual participants can drop in!*
- *Mention availability for today (are you attending talks?)*

General CUPiD Development

Proposer: Mike Levy
Location: Hybrid
Requirements: Laptop / account on NCAR Supercomputers

What: There is a huge backlog of hackathon-appropriate issues, it would be great to work through them

Why: There are several issues that just require a short amount of time to tackle; teaming up to work on them will also provide a chance to entrain new developers in the CUPiD framework

How: I'll welcome help on any open issues, but think it would be helpful to focus on

- Finalizing changes from the LMWG tutorial (does not have "hackathon" label)
- Updating to jupyter-book v2 (sphinx -> myst transition; #332)
- Cleaning up the numeric-execution-count check (#383)

Multi-instance timeseries capability in CUPiD

Proposer: Molly Wieringa
Location: In person
Requirements: Laptop, account on NCAR Supercomputers

What: Many earth system modeling applications involve running ensembles. Within CESM, ensemble simulations can be configured using a multi-instance option (this is regularly used in ensemble data assimilation applications, for example). These multi-instance simulations occupy significant storage space as standard history file output, which would be reduced by postprocessing history / output files into timeseries files. Timeseries functionality exists in CUPiD but is not enabled for multi-instance cases, which use different naming conventions for output files.

Why: This is a low-demand project– very little new code would need to be created. There is an immediate application for multi-instance postprocessing capabilities, and this project could provide a springboard to further expansion CUPiD diagnostics for ensemble simulations.

How: General order of operations

- Introduce multi-instance example case (configuration and file output / naming conventions)
- Discuss potential design choices for adding multi-instance functionality to CUPiD's timeseries tool
- Try things! Potentially comparatively / in parallel

Unit Testing (Fortran and / or Python)

Proposer: Sam Rabin
Location: Hybrid
Requirements: Laptop and a project you want to add testing to

What: It's often said that untested code is broken code, but too often we as software engineers are not able to prioritize adding tests. To some extent, this may stem from not being familiar enough with the available testing tools—but the more we use them, the lower the effort. In this hackathon, I will briefly introduce two testing frameworks: `pytest` for Python and `pfUnit` for Fortran. We will then dedicate the rest of our time to adding tests to our projects.

Why: Learning the basics of these frameworks will take hardly any time at all. Adding tests is then not a large project—participants will be able to add as many tests as they have time for, with even just one simple test representing a keepable product. Even people already familiar with these frameworks may find it useful to have dedicated time for adding tests, rather than trying to fit them in during general development work.

How: I will present for 15-30 minutes on each of `pytest` and/or `pfUnit`, depending on participant interest. The rest of the time will be dedicated to getting participants set up with the frameworks and then adding tests. I would welcome others familiar with these tools to attend as well, to help troubleshoot.

Using the AMReX a C++ productivity framework to modernize MOM6

Proposer: John Dennis
Location: hybrid
Requirements: Laptop / NCAR supercomputers

What: The aim of the TURBO project is to modernize MOM6 such that it supports execution on both CPU and GPU systems. This will be achieved by leveraging AMReX a C++ productivity framework to create two distinctive software products. The first is an alternative to the existing Flexible Model System (FMS) on which MOM6 is currently based, while the second is a prototype of a what MOM might look like if it was written using the AMReX framework.

Why: Currently there are a lot of small units of work that can be easily divided amount different developers.

How: We are currently working on writing unit tests for 53 critical interfaces that MOM6 uses to call into FMS using pfUnit. We are also implementing each of these 53 critical interfaces using AMReX.

Using an AI Assistant

Proposer: TBD (Brian D?)
Location: Hybrid
Requirements: Laptop

What: Someone who knows what they are doing will walk through an introduction to coding with an AI Assistant, and then participants can work on small projects together or individually to get familiar with the tool

Why: After hearing the AI Assistant talks in the Thursday morning session, it could be useful to get hands-on experience

How: The most common setup will probably be using Claude via VS Code? [Mike is writing this up and he doesn't know anything]

Transition to Morning Session

AI Assistants

9:30 - 9:45 Sam Rabin (NSF NCAR / CGD)
9:45 - 10:00 David Appelhans (NVIDIA)
10:00 - 10:15 John Dennis (NSF NCAR / CISL)
10:15 - 10:30 Michael Kavulich (NSF NCAR / RAL)
10:30 - 11:00 BREAK
11:00 - 11:15 Aya Lahlou (Columbia; remote)
11:15 - 11:20 Kate Thayer-Calder (NSF NCAR / CGD)
11:20 - 11:45 Discussion

Hackathon

If you want to work on projects, please step into Outer Damon Room and self-organize