The Community Earth System Model

The CESM is *critical computational infrastructure* that enables a large Community of scientists to study a diverse set of Earth-system processes across a broad range of time and space scales.

The Community is not monolithic. It is diverse.

Priorities have to be set, but important segments of the Community are currently under-served, and scientific opportunities are being missed.

Left behind

For the last 20 years, modeling centers around the world have been developing very high-resolution global models *for both weather and climate research*.

The work has now reached a level of maturity. It is no longer exotic. It is mainstream.

CESM is not in the game.

We're here to help.

CLIMATE CHANGE

Europe builds 'digital twin' of Earth to hone climate forecasts

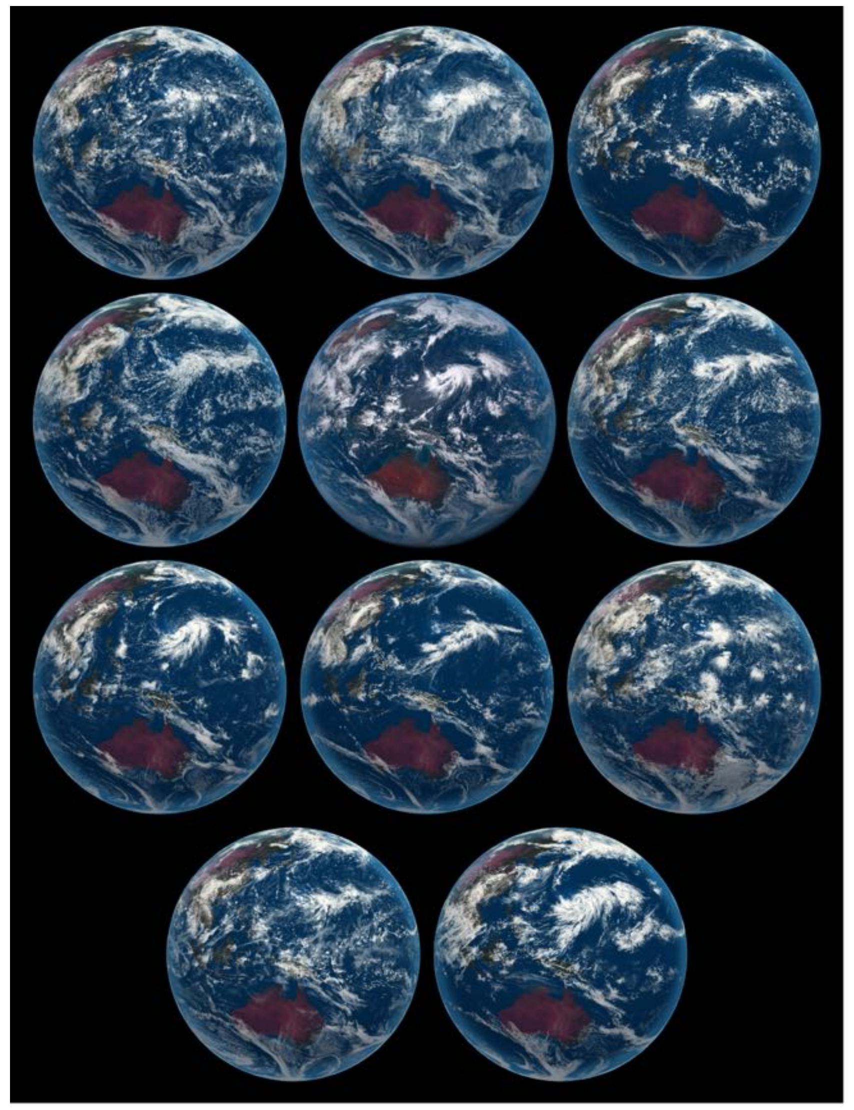

DYAMOND

Proposal to the National Science Foundation's
Program for Cyberinfrastructure for Sustained Scientific Innovation for

Collaborative Research: Frameworks:

Community-Based Weather and Climate Simulation With a Global Storm-Resolving Model

Principal Investigator:
David A. Randall, Colorado State University

Co-Principal Investigators:
James Hurrell, Colorado State University

Andrew Gettelman, National Center for Atmospheric Research

Richard Loft, National Center for Atmospheric Research

William Skamarock, National Center for Atmospheric Research

Submitted November 1, 2019

EarthWorks

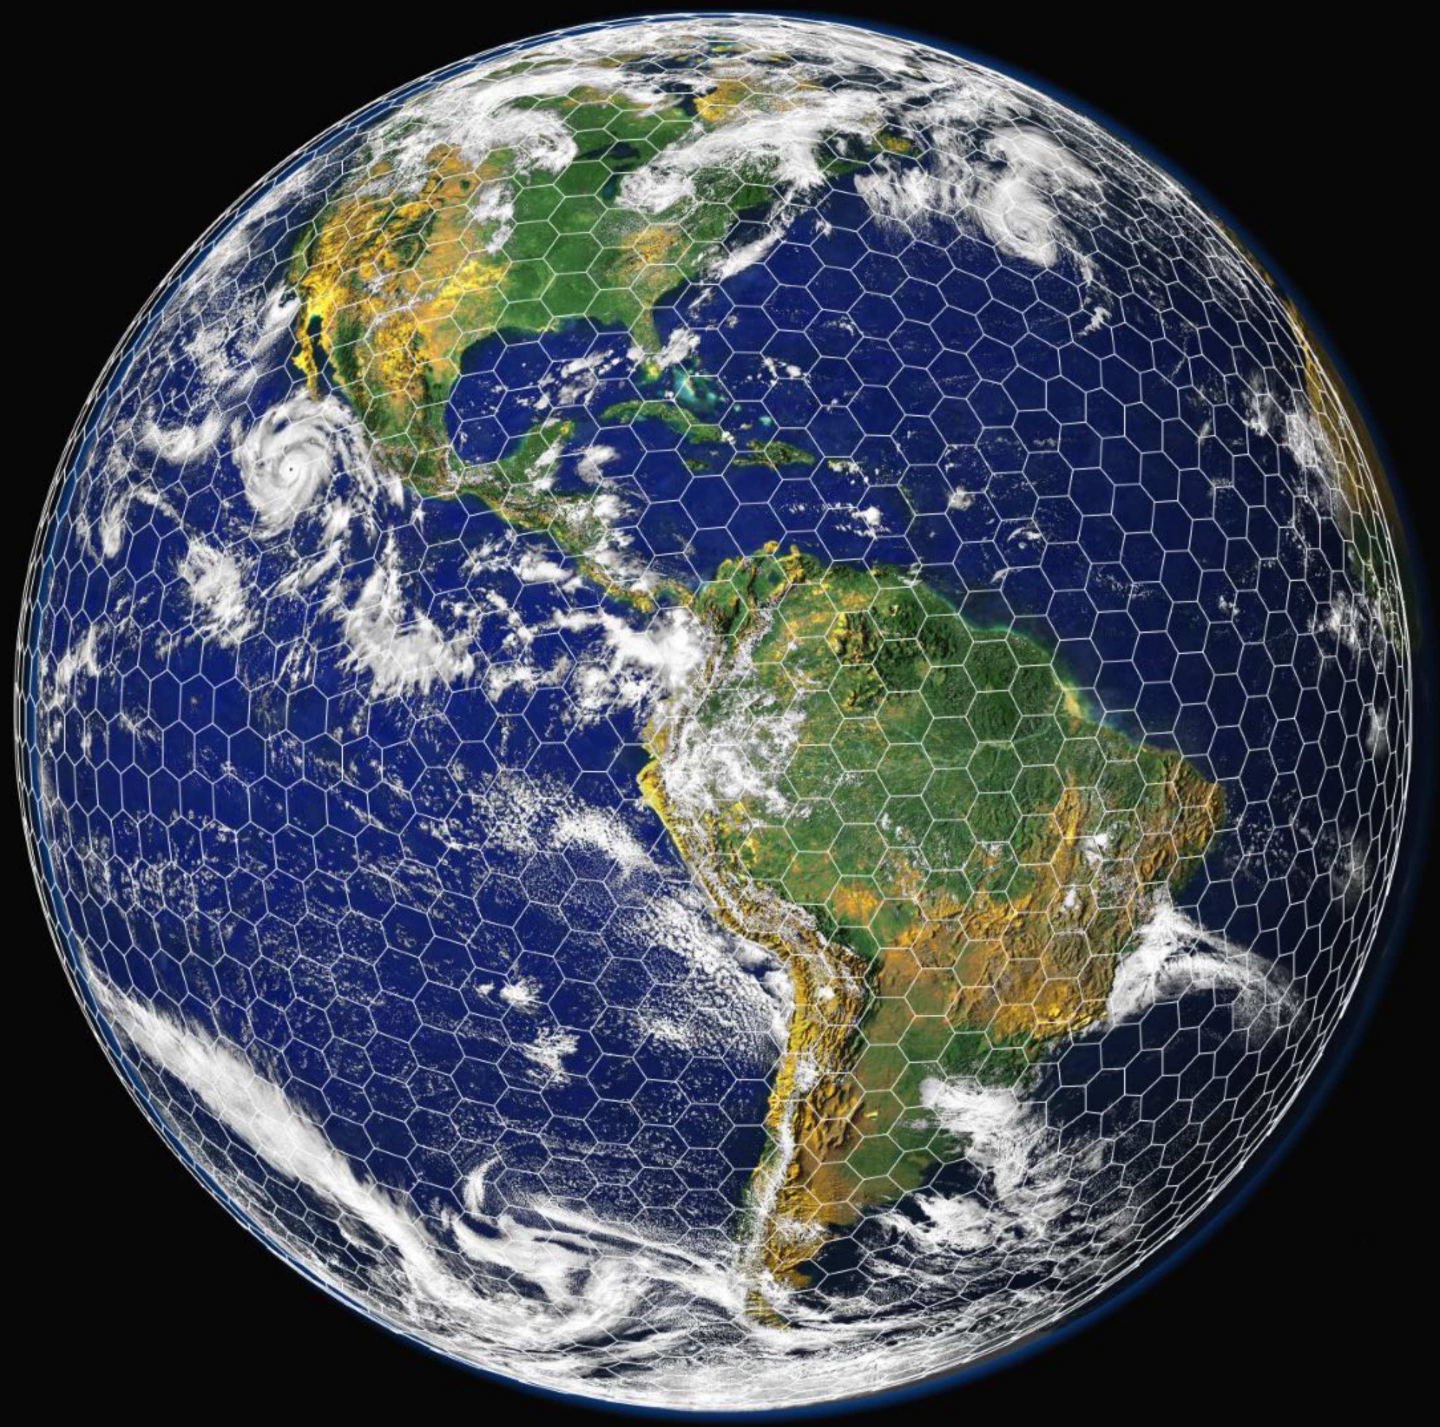

EarthWorks is led by Colorado State University, with Co-PIs from three NCAR laboratories.

CSU

CSU

MMM

CGD

CISL

EarthWorks

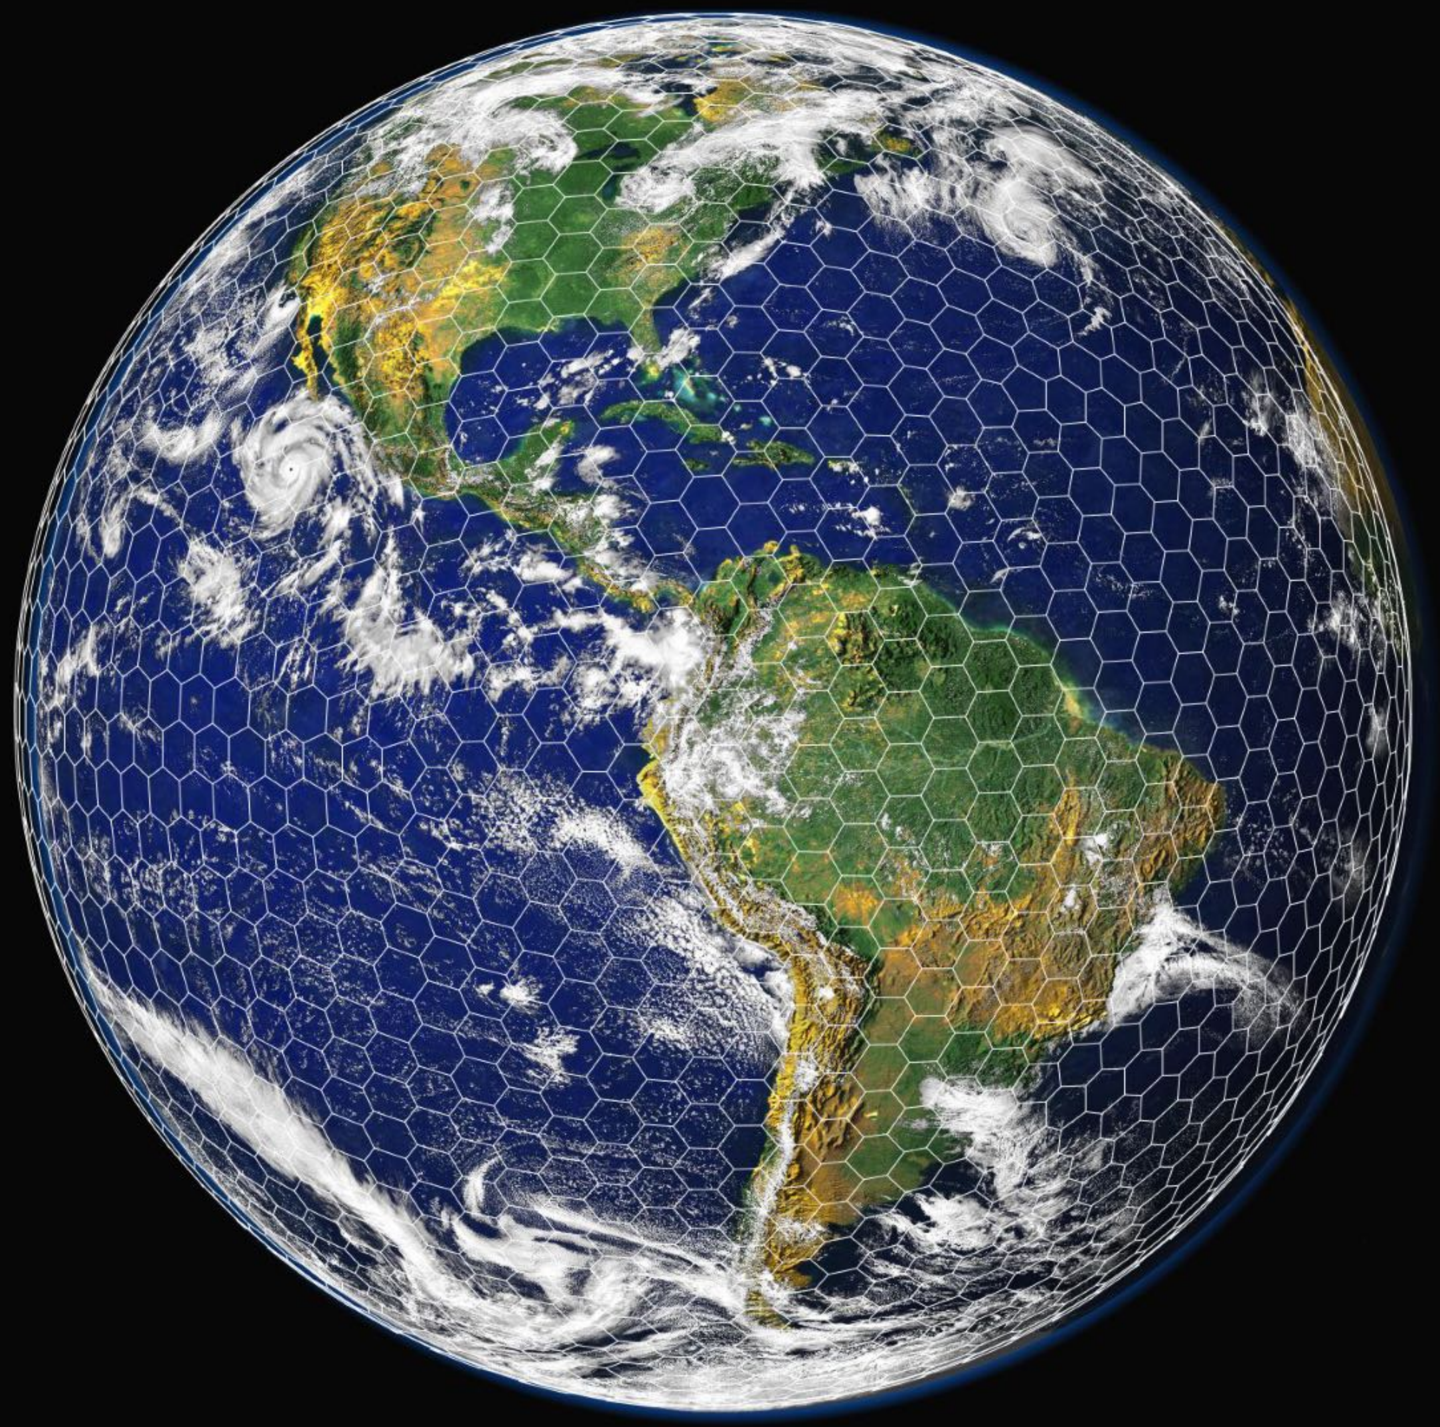

EarthWorks is a five-year project to develop a global coupled model, based on the CESM, that uses a *~4 km global grid* for the atmosphere, ocean, and land surface.

The model will be used to study both weather and climate.

CSU

CSU

MMM

CGD

CISL

What a ~4 km grid can do

For the atmosphere and land surface:

- Thunderstorms & mesoscale convective systems
- Hurricanes of realistic intensity
- Individual large mountains and valleys
- Gravity waves
- Coastlines
- Many lakes, and large rivers
- Cities

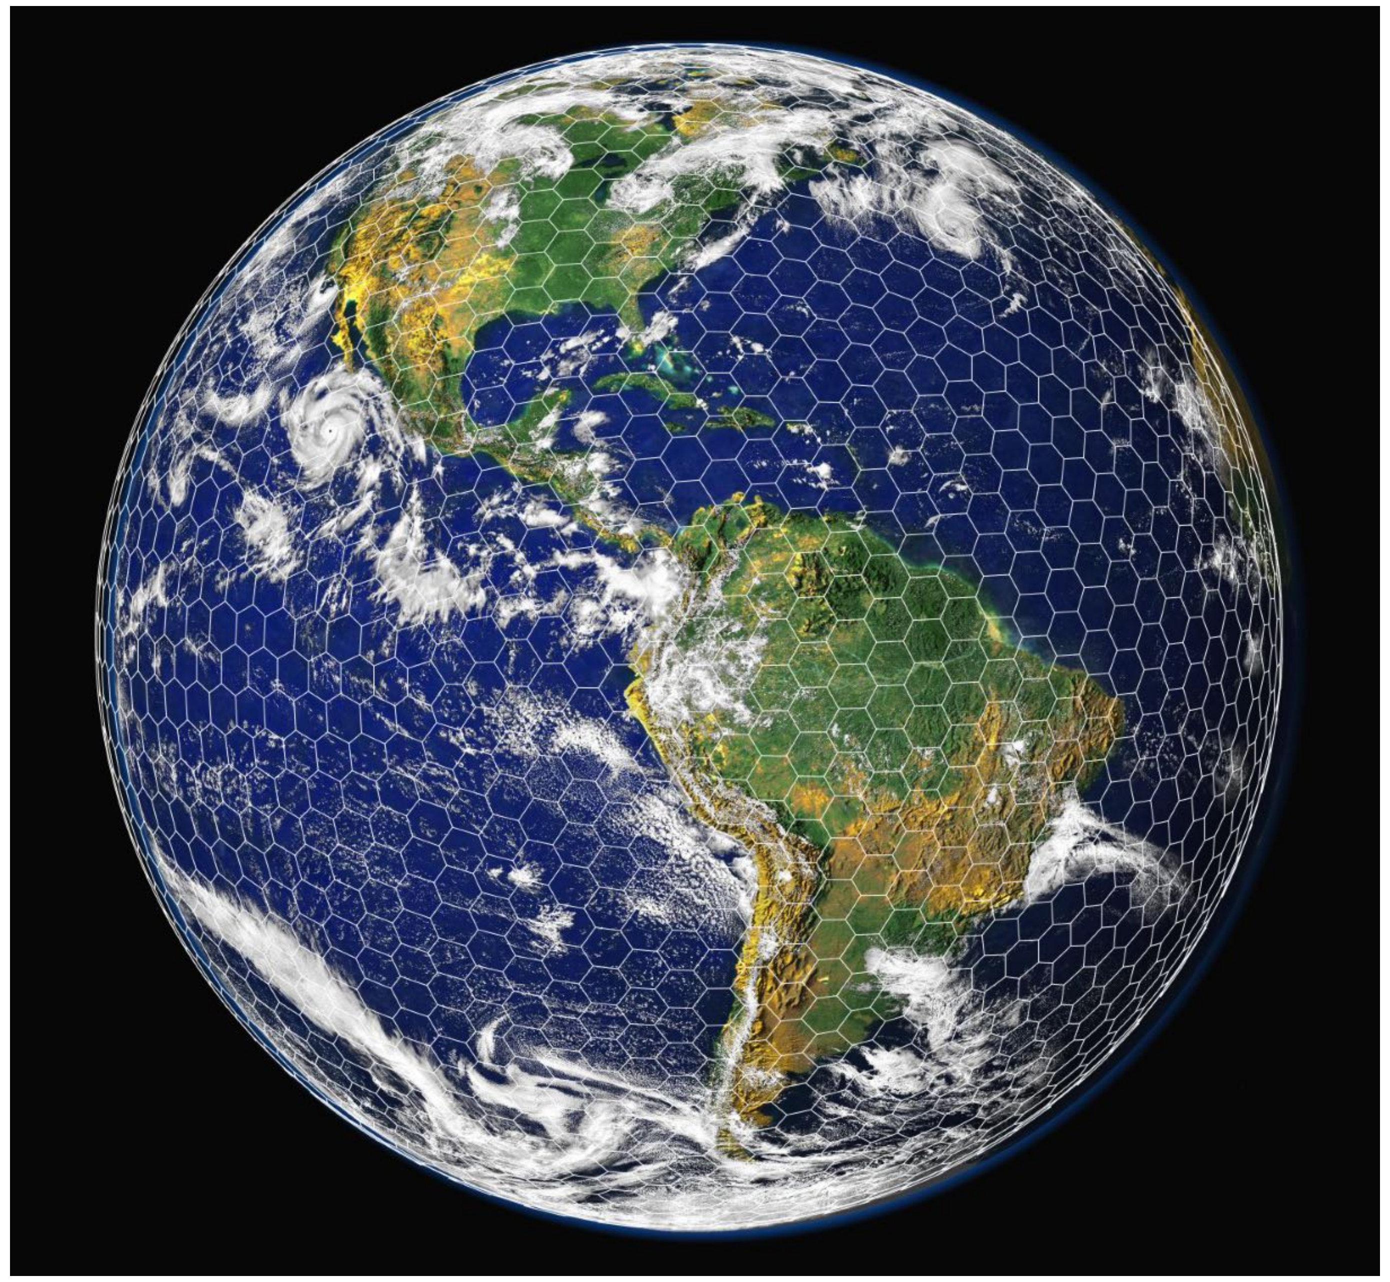

For the ocean:

- The most energetic eddies
- Deep convection
- Bottom topography
- Gravity waves
- Estuaries

Changes

With a 4 km grid, the design of the model has to change.

Change is an opportunity, not a problem.

- Non-hydrostatic dynamics
- Optimal grids
- Partially explicit simulation of deep convection for the atmosphere and ocean
- Explicit gravity-wave drag for the atmosphere and ocean

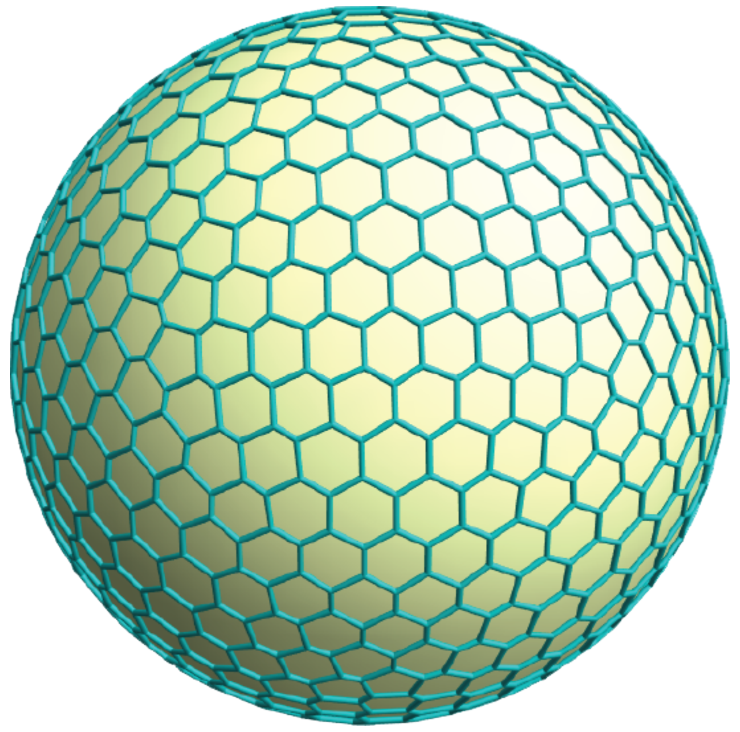

EarthWorks components

- The atmosphere model: A modified version of CAM6
  - MPAS atmosphere non-hydrostatic dynamical core, developed by MMM
  - A resolved stratosphere
  - High-resolution CAM-ish physics
  - The Community Physics Framework (CPF, a.k.a. CCPP)
- The MPAS ocean model developed at Los Alamos and used by E3SM
- The MPAS sea ice model, which is based on CISE and designed to work with the MPAS ocean model
- The Community Land Model (CLM)
- The Community Mediator for Earth Prediction Systems (CMEPS)

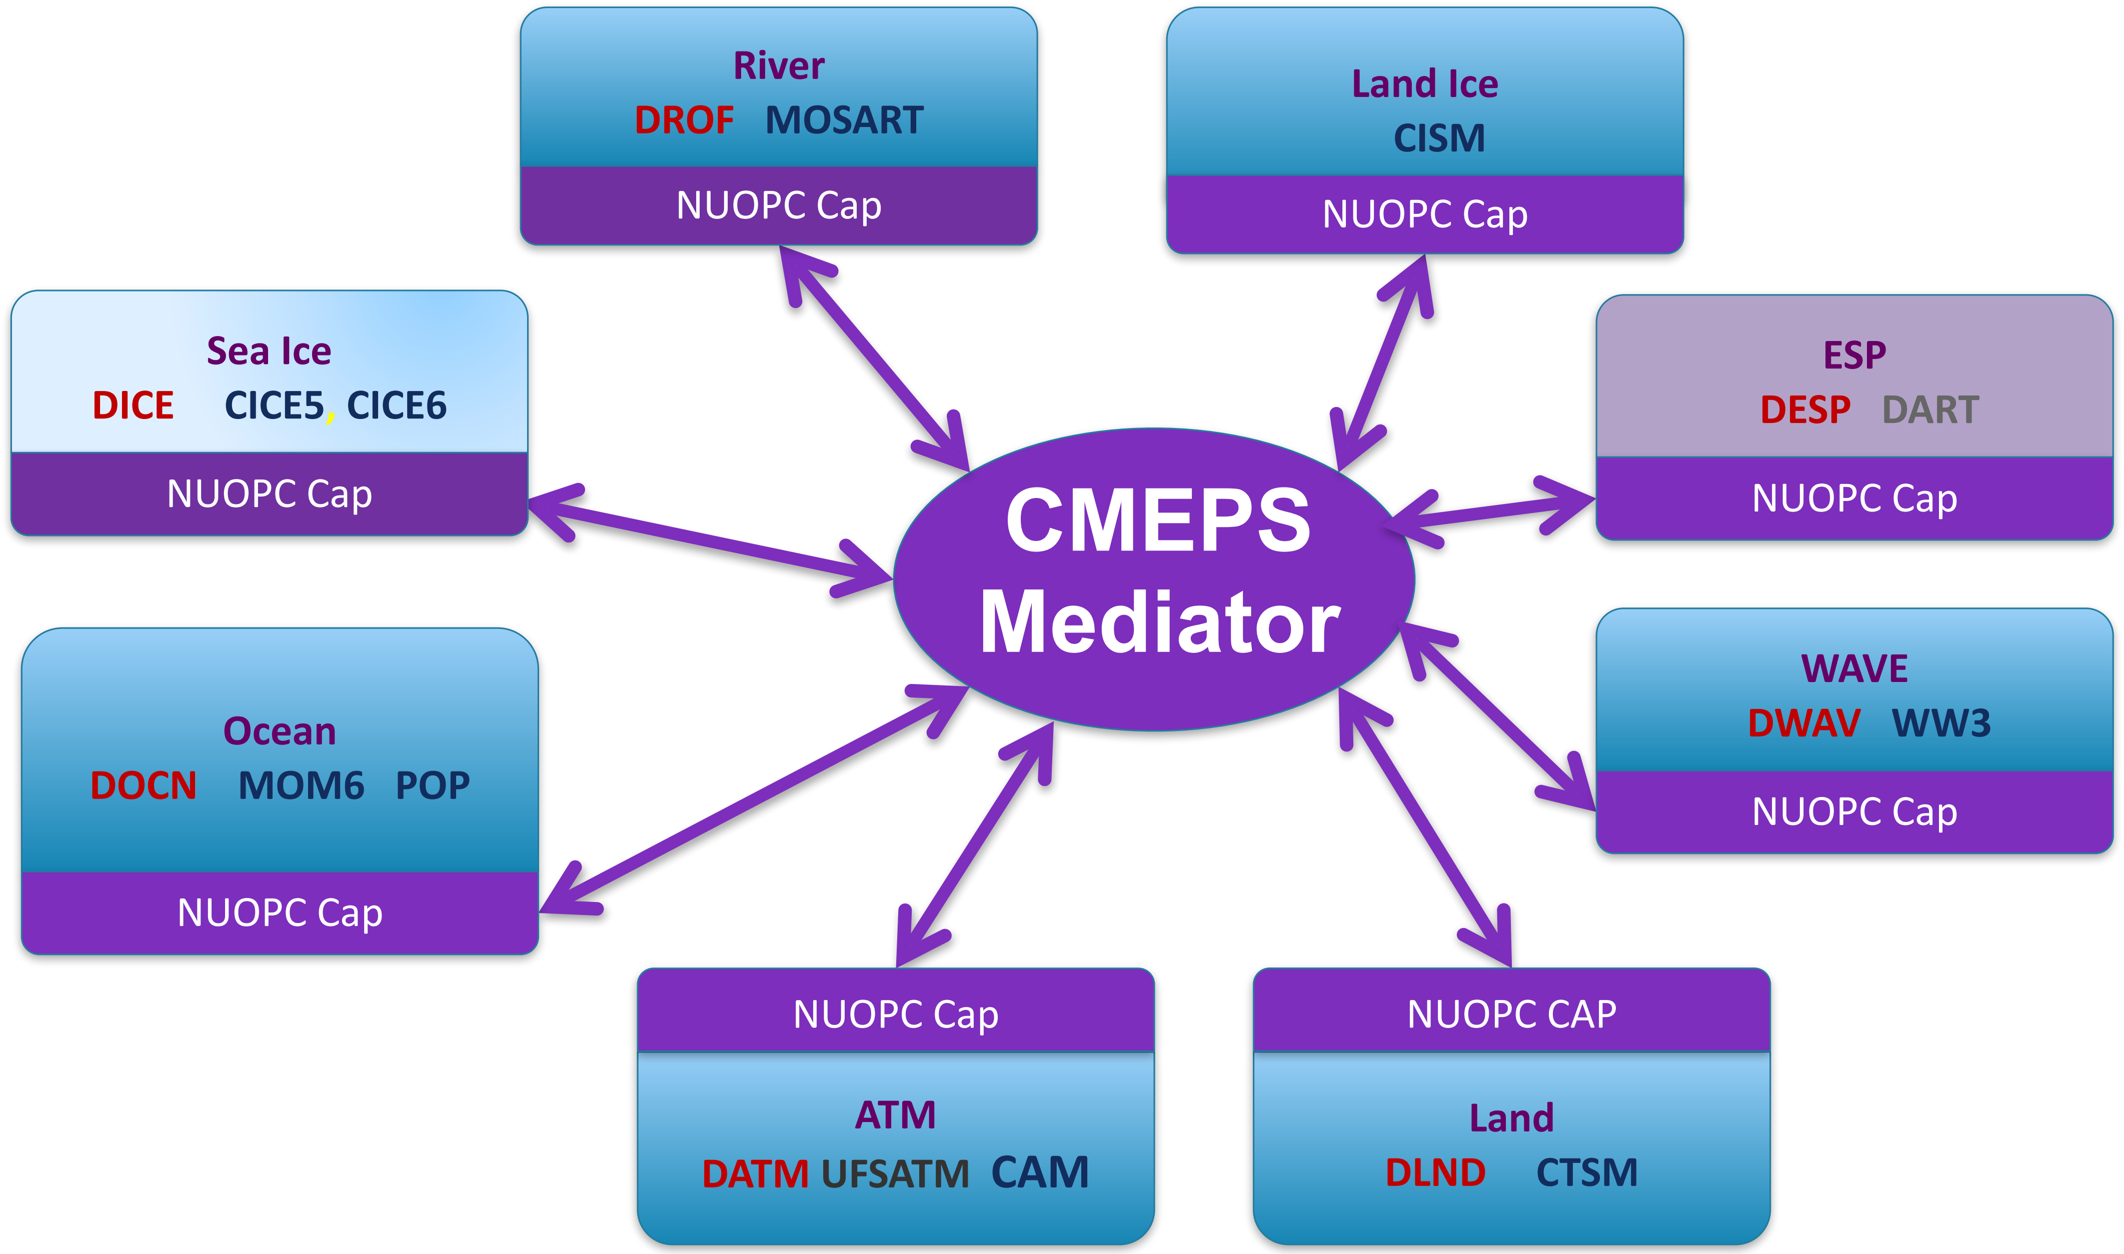

EarthWorks will not do this.

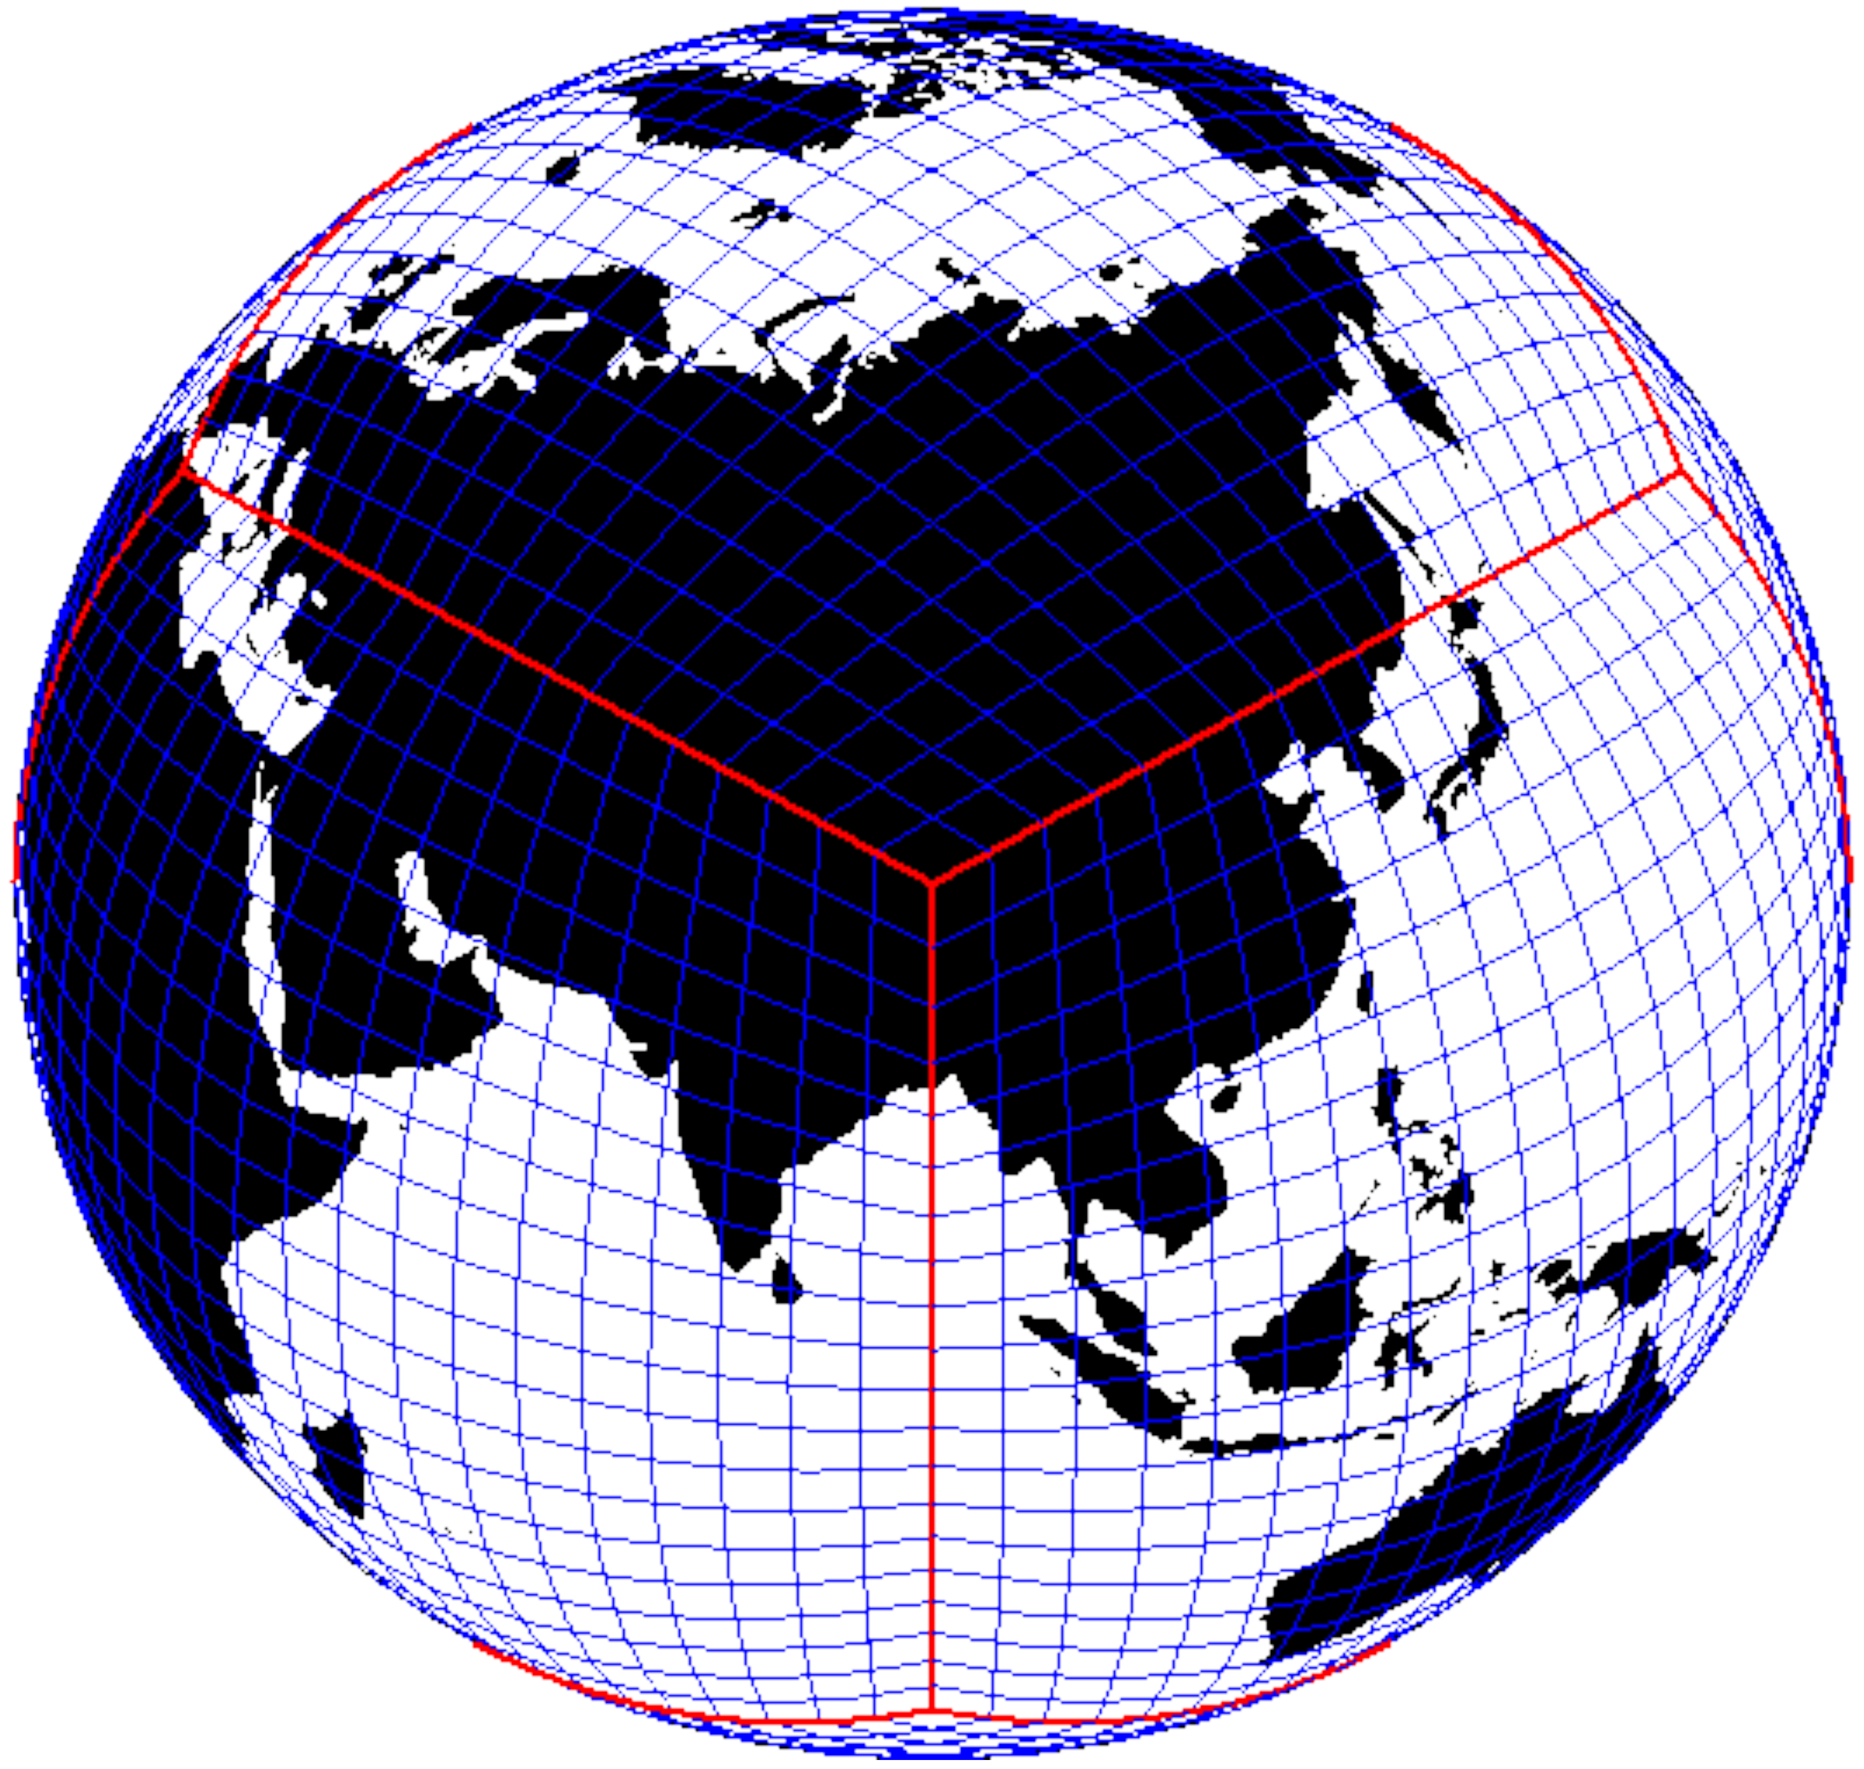

Cubed sphere atmosphere grid

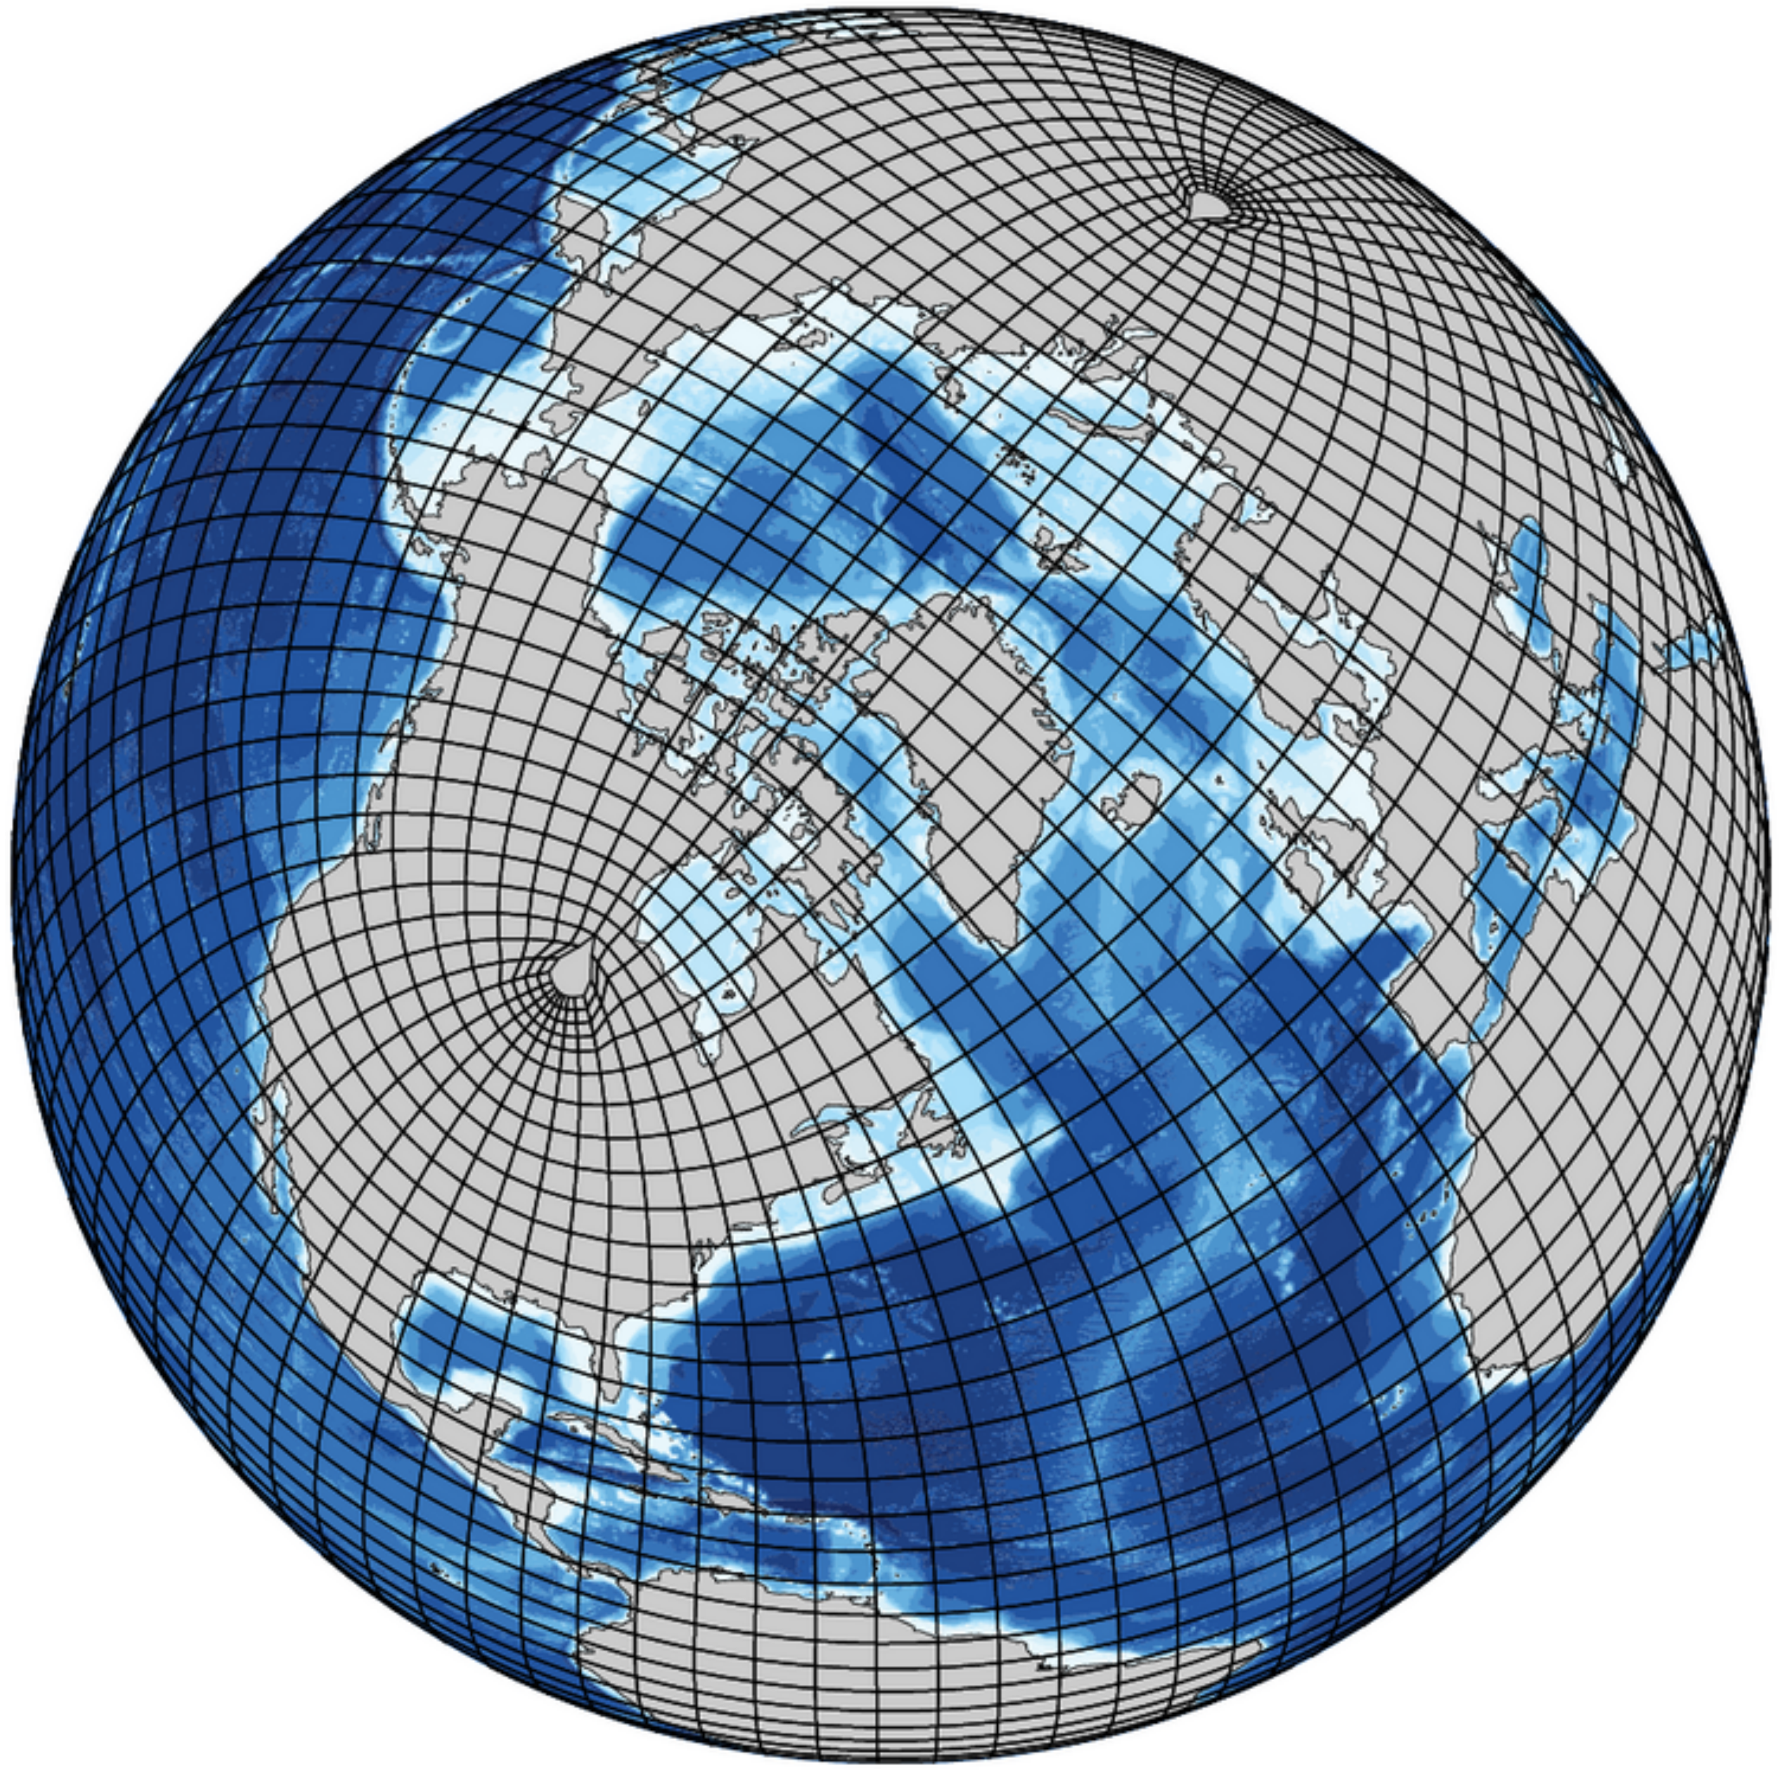

Tri-polar ocean grid

One grid to rule them all

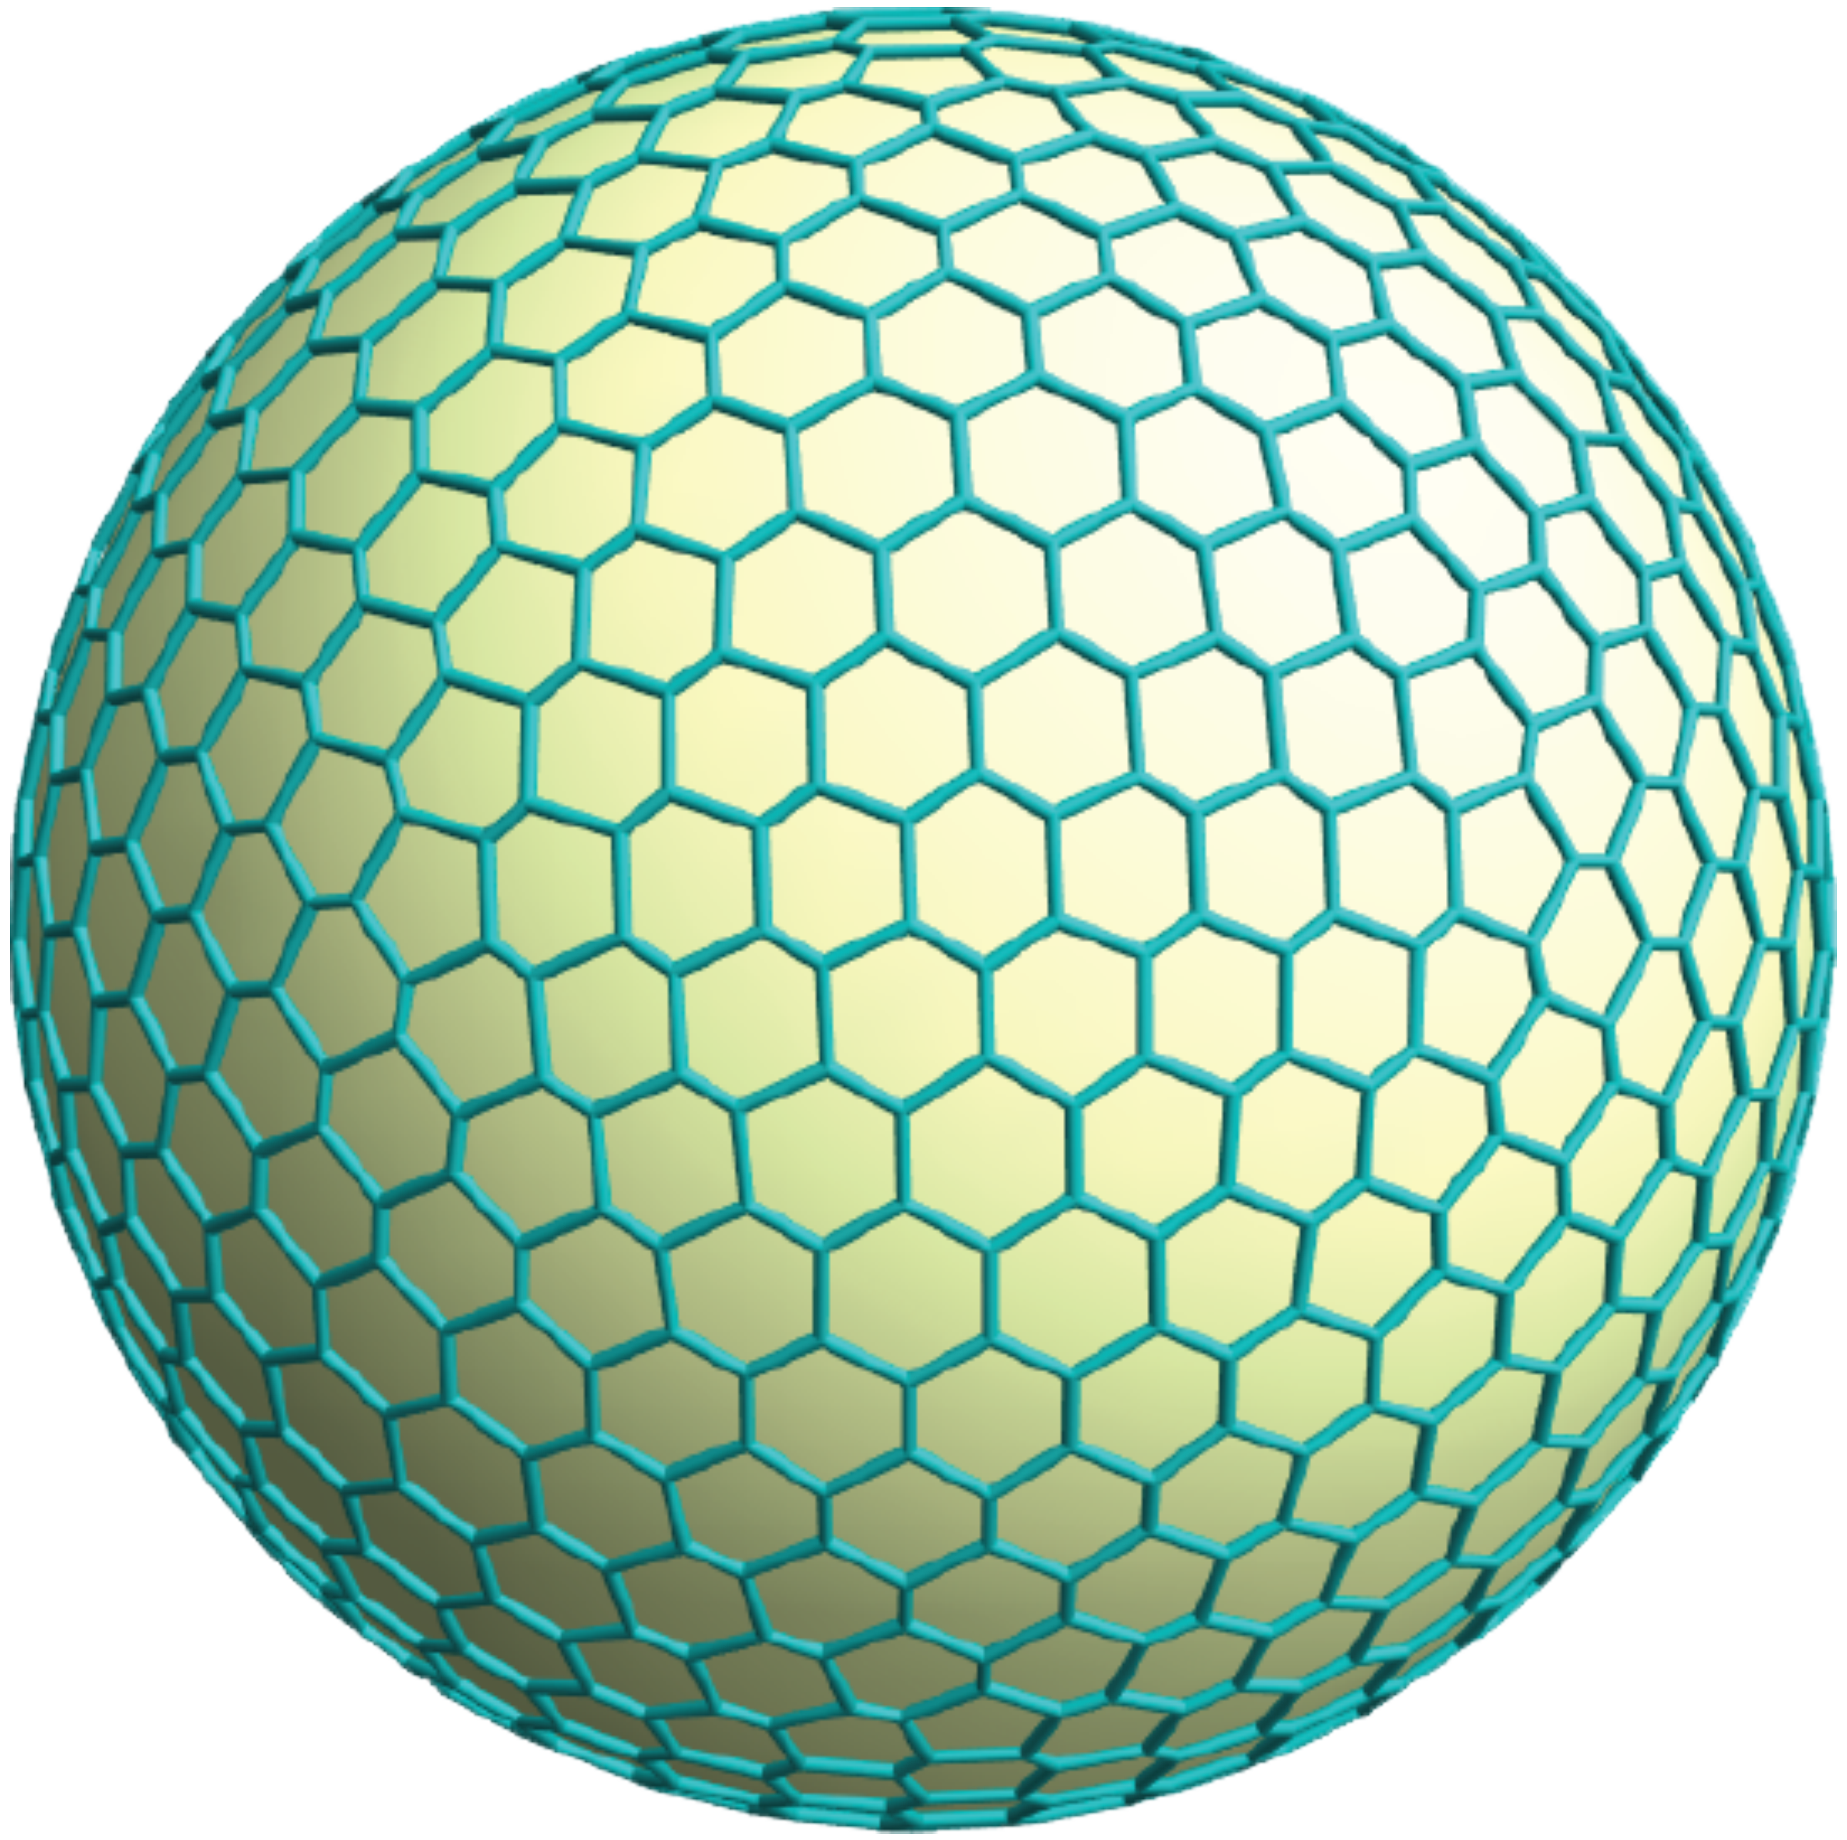

Grid	No. of grid points N	Avg grid distance ℓ (km)
G0	12	6699.1
G1	42	3709.8
G2	162	1908.8
G3	642	961.4
G4	2562	481.6
G5	10 242	240.9
G6	40 962	120.4
G7	163 842	60.2
G8	655 362	30.1
G9	2 621 442	15.0
G10	10 485 762	7.53
G11	41 943 042	3.76
G12	167 772 162	1.88
G13	671 088 642	0.94

Non-hydrostatic regime

Target grid spacing

Simplify

As discussed above, a grid spacing of ~4 km works well for the atmosphere, the ocean, and the land surface.

The single-grid architecture of *EarthWorks,* combined with CMEPS, will enable a lower operation count, less message-passing overhead, and reduced memory requirements.

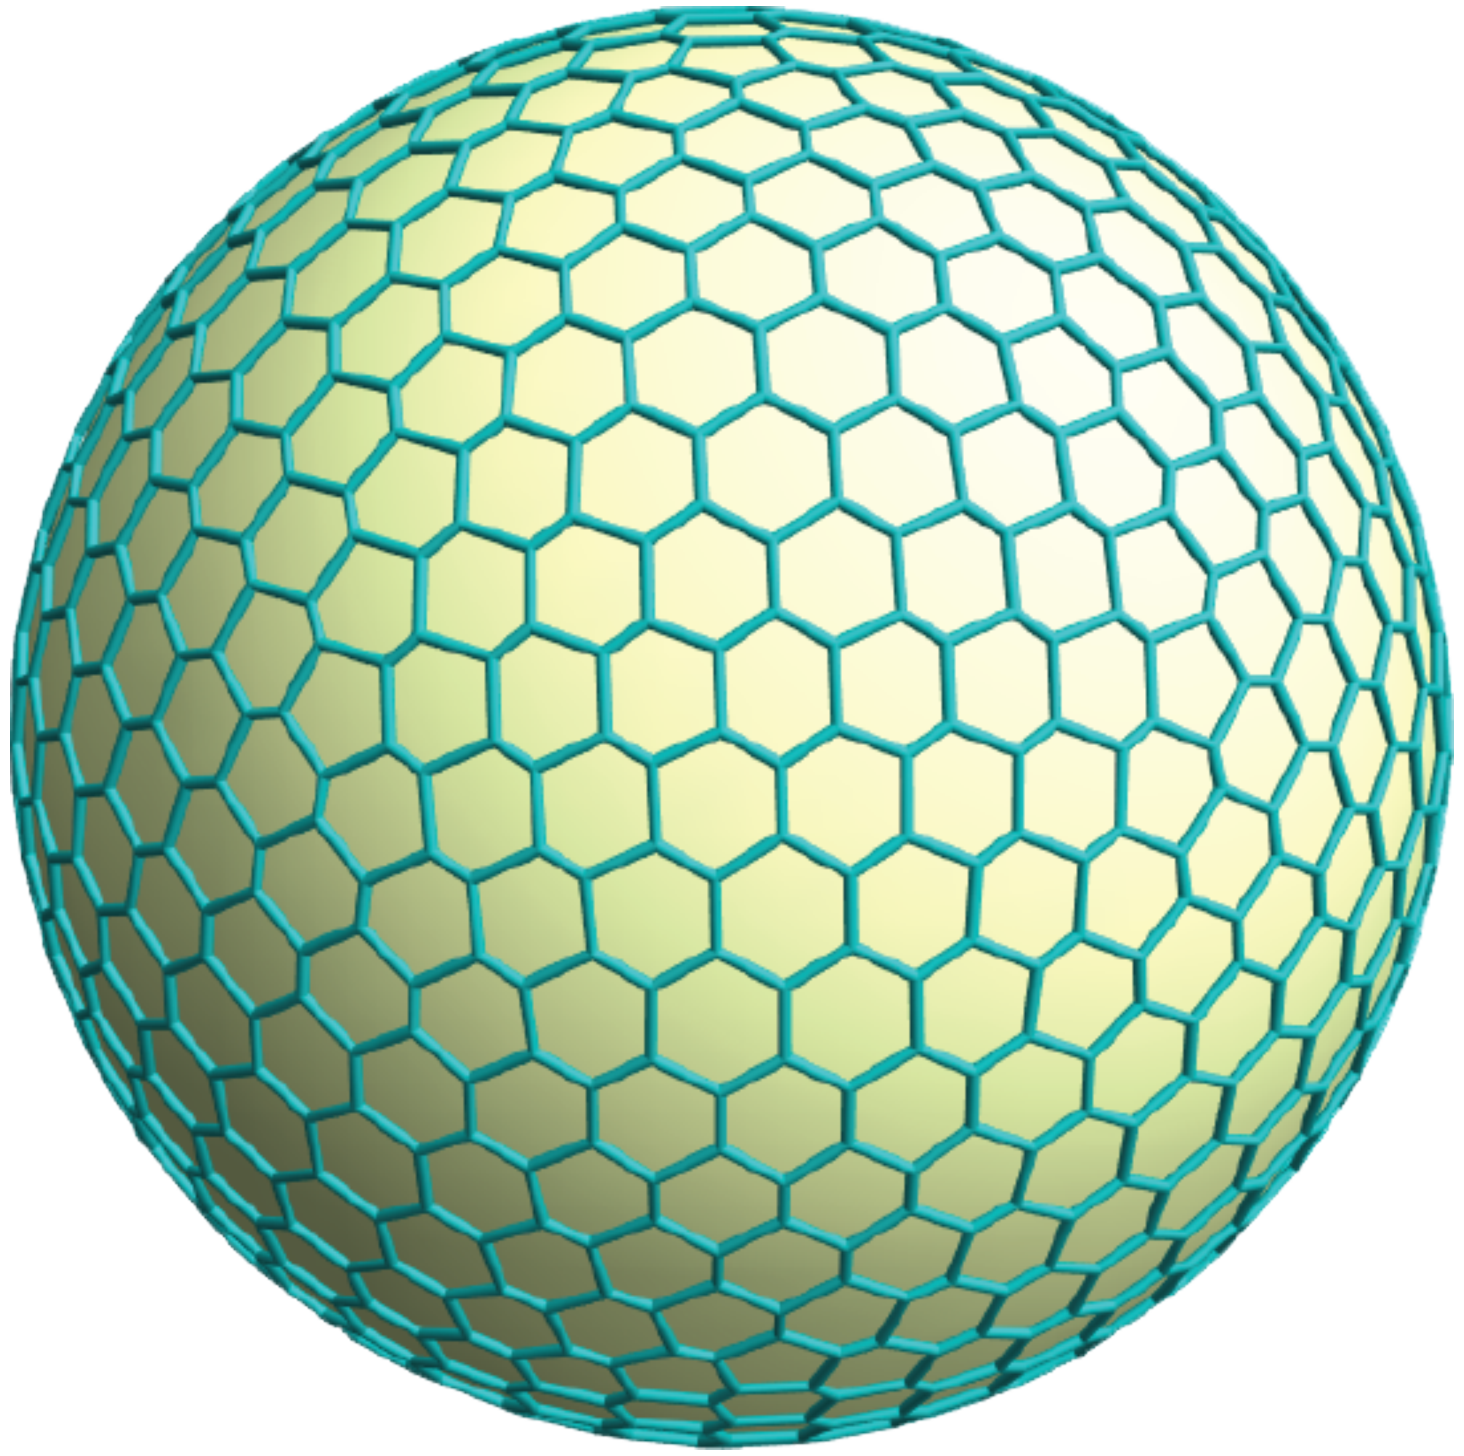

With cloud-permitting resolution for the atmosphere and eddy-resolving resolution for the ocean, a single grid is the right way to go.

Science Goals

Simulation of
extreme events
across scales

Understanding the
resolution-dependence of
convection and turbulence

Understanding gravity
wave generation and
propagation

Simulation of global
climate impacts on
human scales

Performance goals

Our 2025 performance goals for a version of EarthWorks with ~4 km global grid spacing are:

- Half a simulated year per day in atmosphere-only simulations with high vertical resolution, and
- One simulated year per day in coupled simulations with fewer layers.

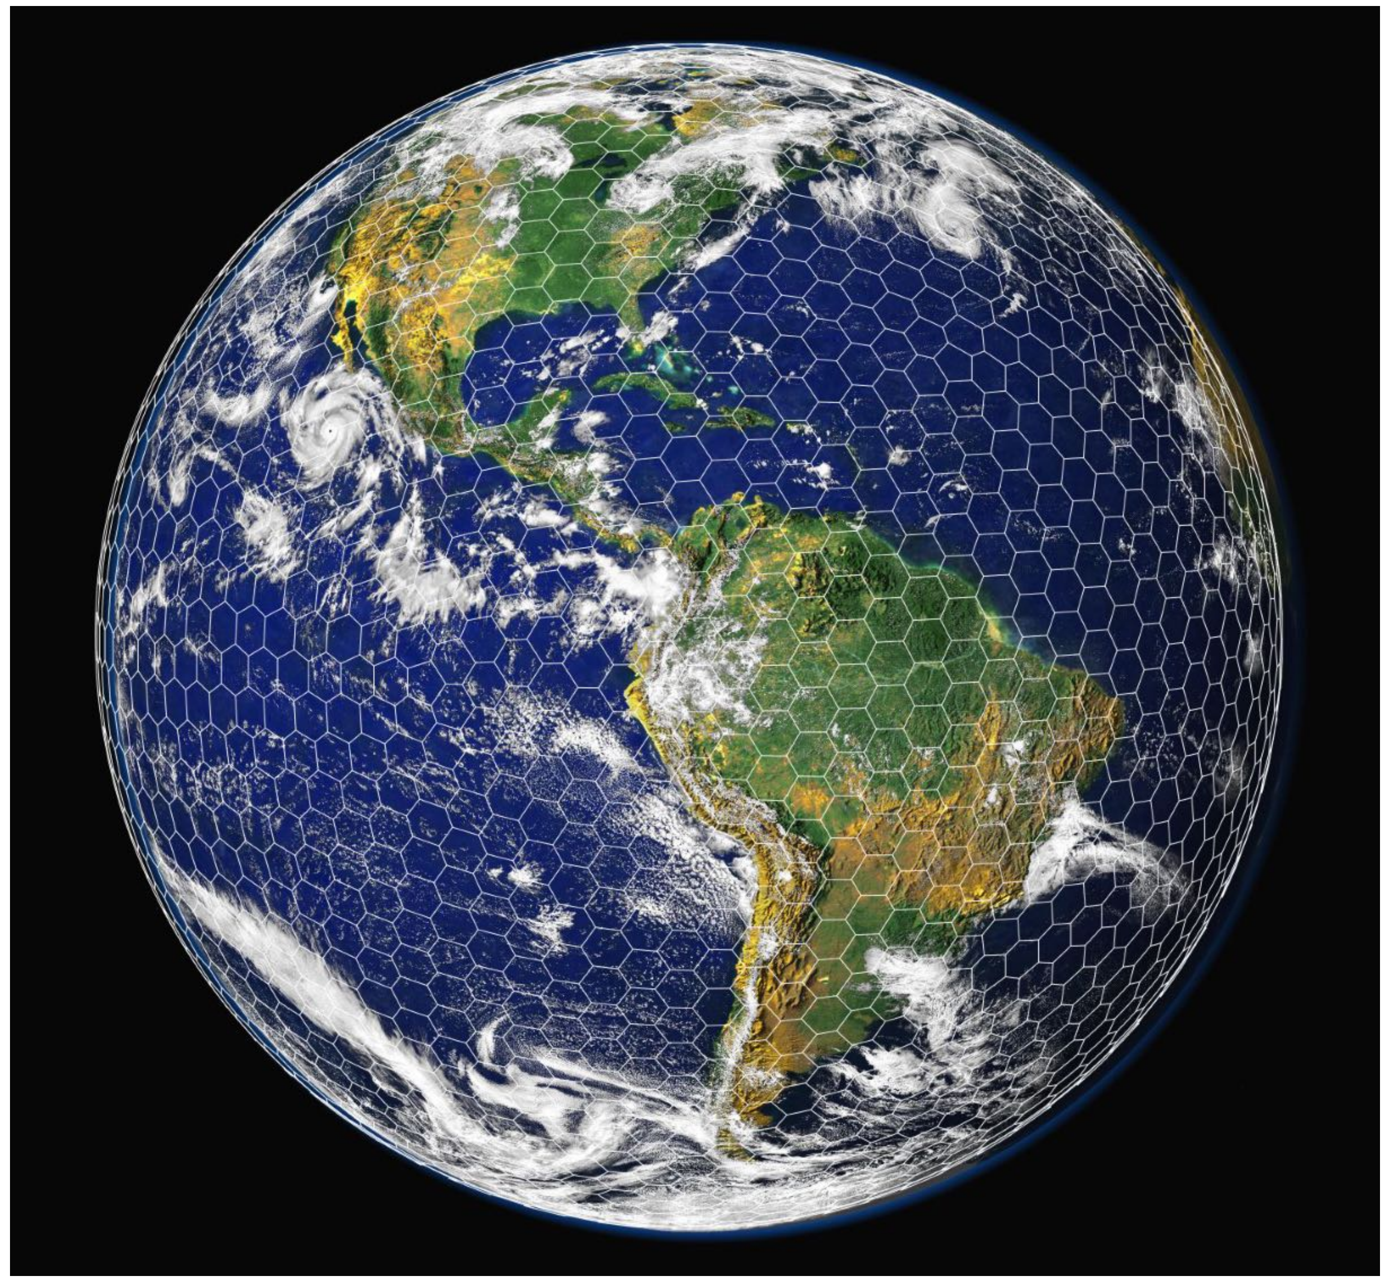

Is 0.5 SYPD at 4 km in 2025 feasible?

- Short answer: Yes, with caveats
- We have to extrapolate from moist MPAS tests
  - 10 km shows 1.64 SYPD on 324 GPUs.
  - Adjust to 4 km climate timestep (2.5x - CFL) ;
  - Adjust to 100 levels (1.8x - observed)
  - Physics overhead (2x - measured);
  - SP->DP (1.5x- measured)
- Answer: 0.12 SYPD
- Cost of additional tracers (6->??) not included.
- Can 4x come from better GPU hardware?
- Close. Both GPU floating point and memory performance is improving ~1.5x every 2 years, and we're due for a refresh this year.

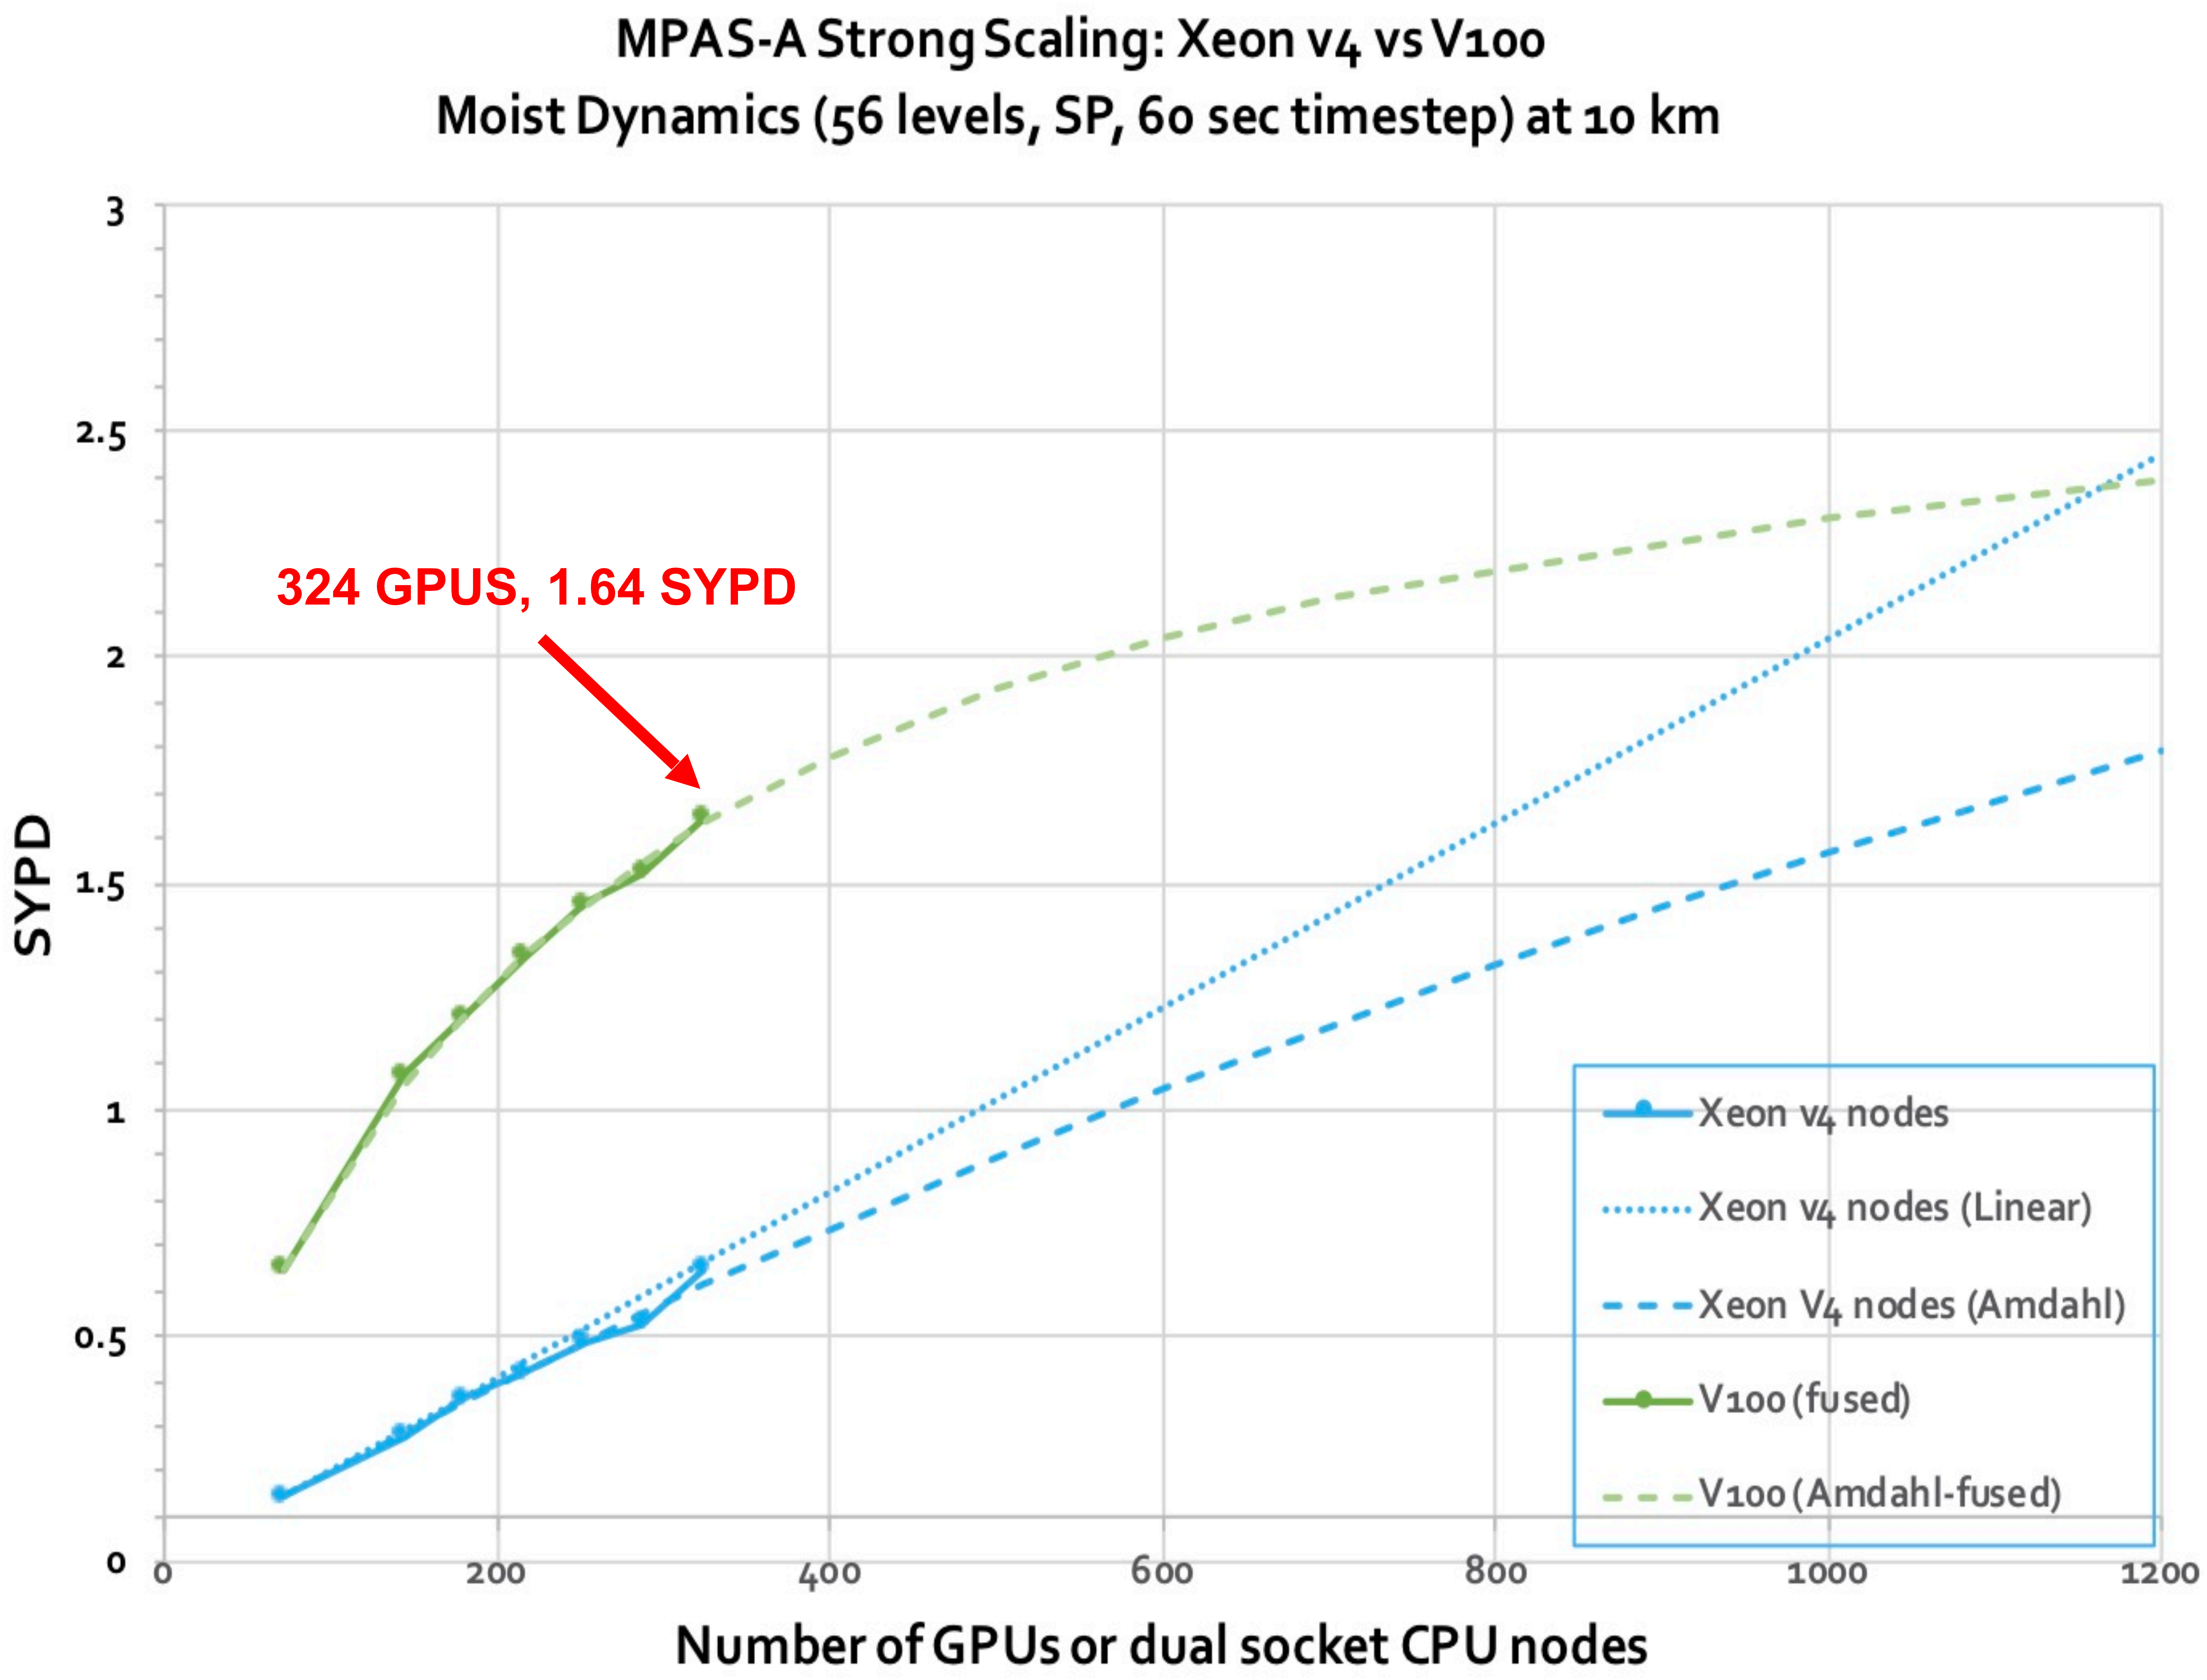

Comparison of MPAS-O (EC60to30 CASE) and MOM6 (T62_t061) on Cheyenne

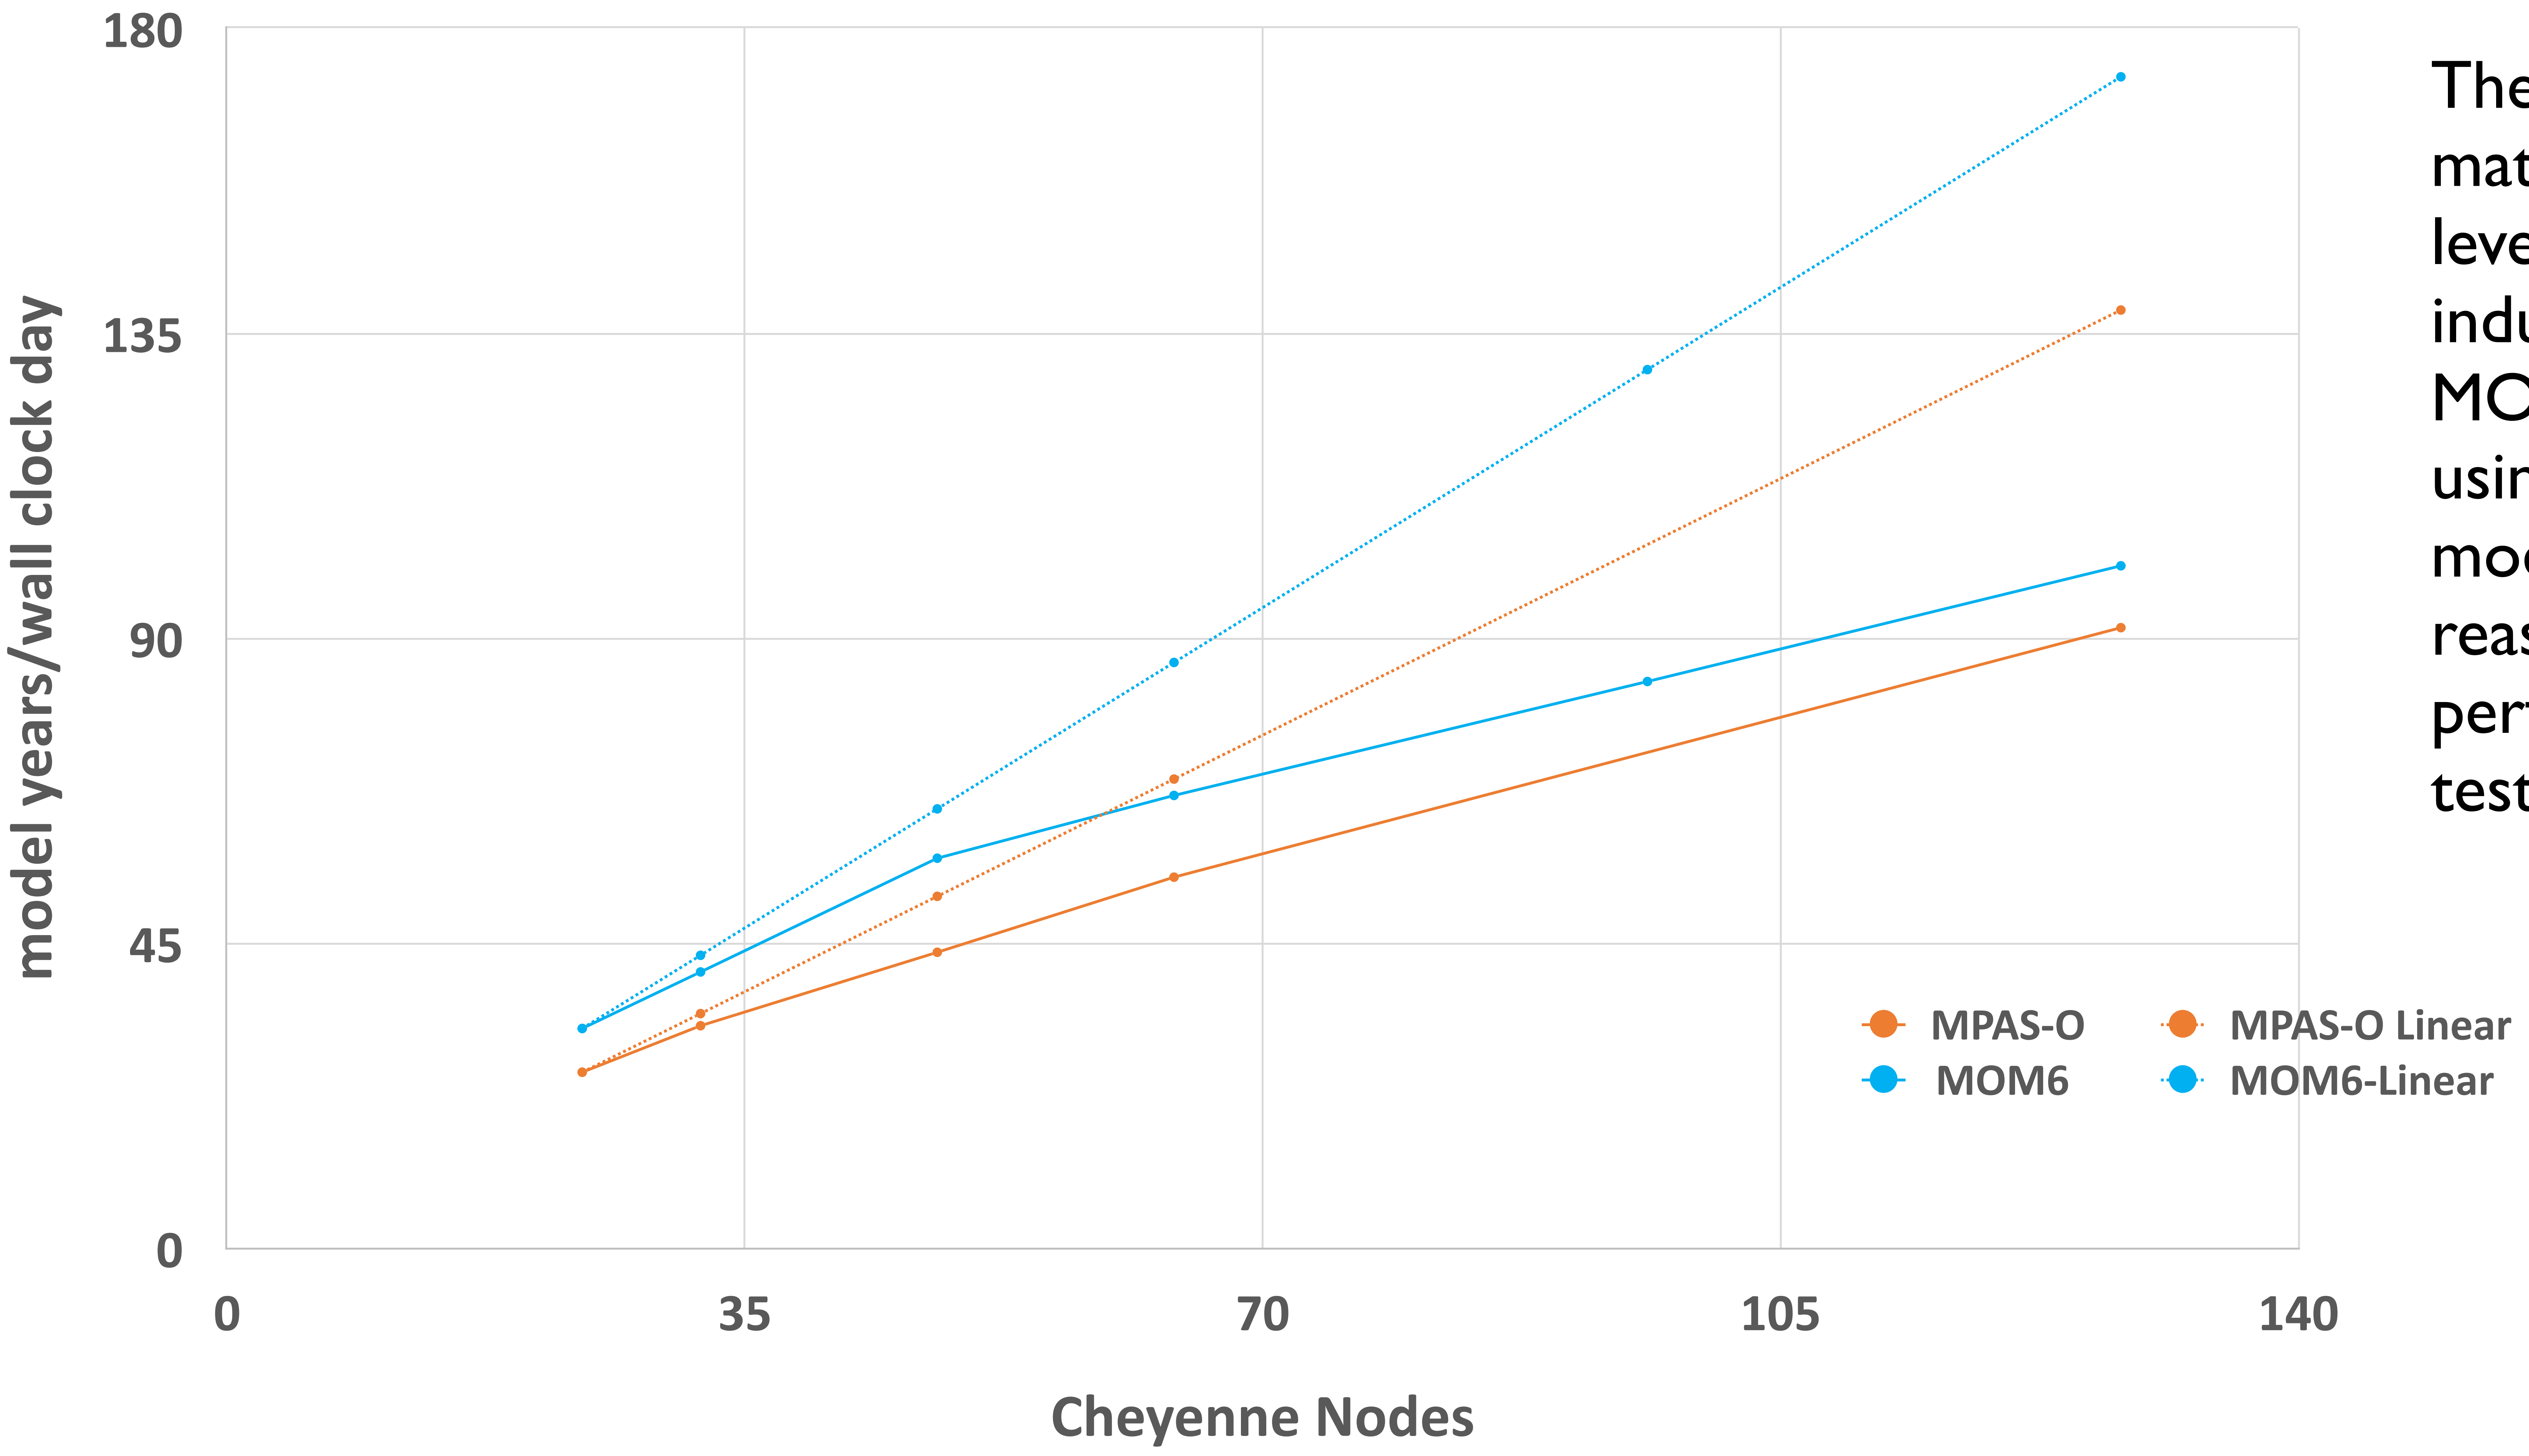

The two ocean model cases are well matched in resolution and number of levels, allowing comparison without inducing excessive scaling errors. The MOM-6 example is run within CESM using the C compset (ocean + data models), allowing us to make a reasonable estimate of the performance of MPAS-O in a similar test configuration.

"Unlike many other high-resolution models under development or production that are atmosphere only, EarthWorks' target to develop a coupled model has potential to answer many questions that other models in the same class cannot."   — Panelist Comment

Software development

- **GPUs**
  - Refactoring for performance portability through directives.
  - Leveraging exascale systems being deployed or contemplated.
- **AI components**
  - Fortran-compatible inference engine
- **Big Data Tools**
  - Parallel post processing
  - Data compression tools
  - Visualization tools — John Clyne’s project

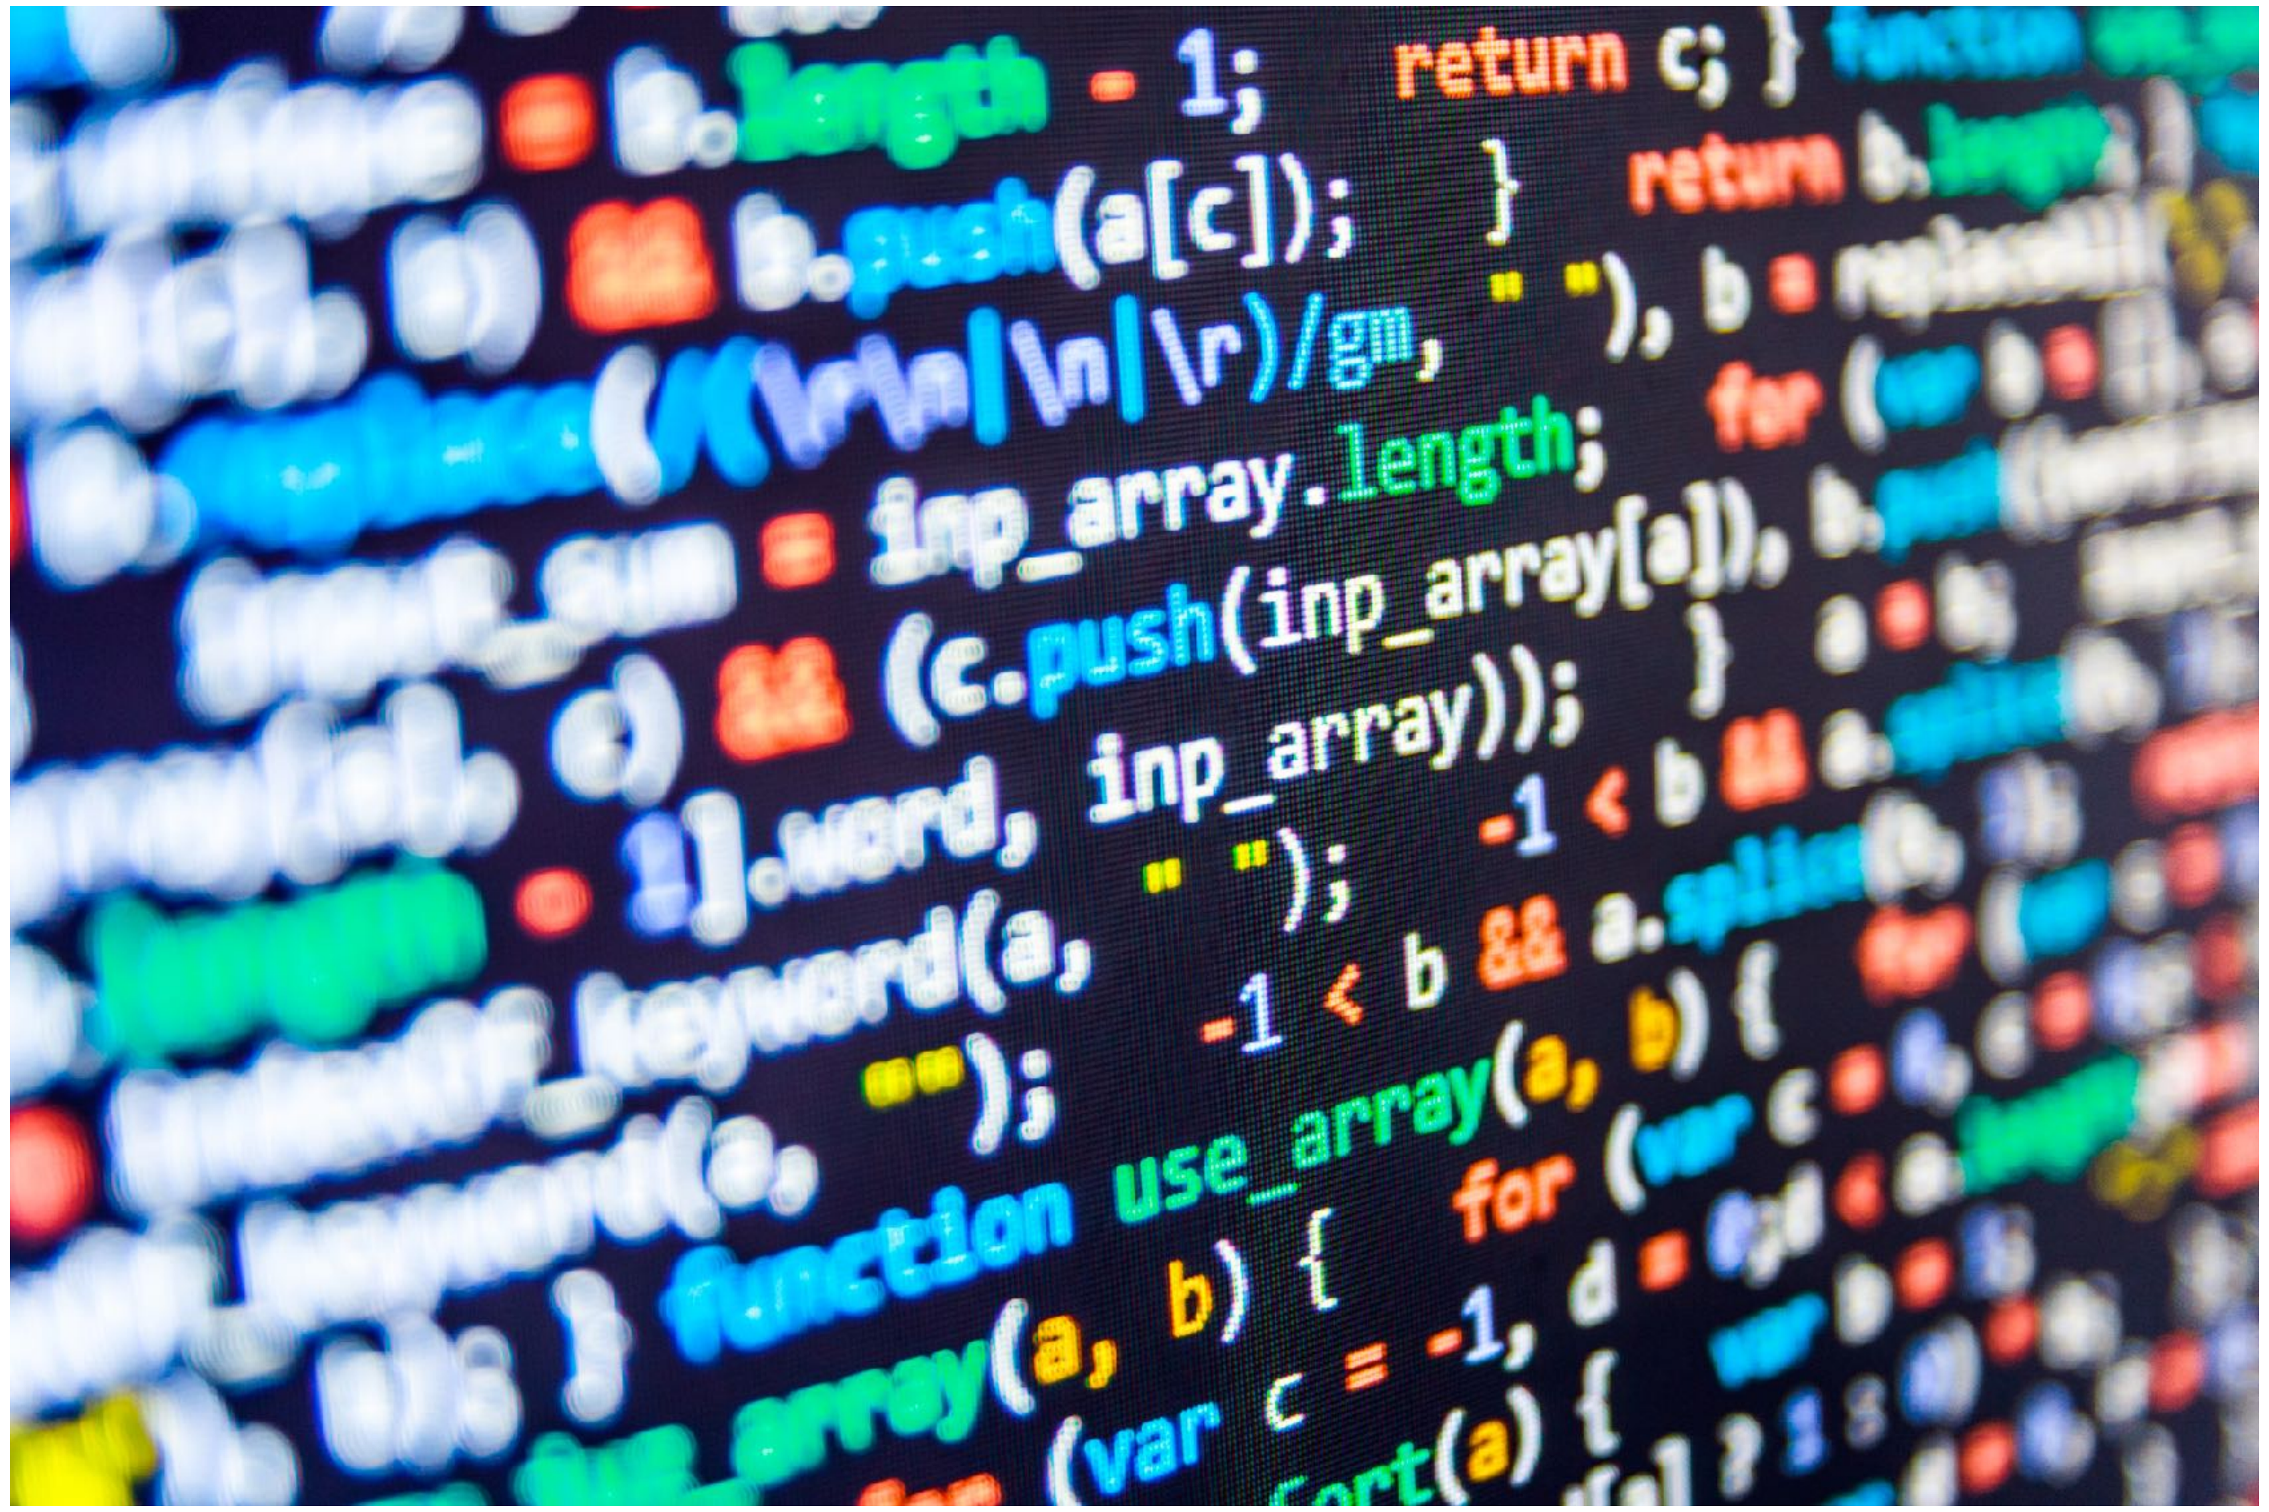

Year 1: integration, debugging, and testing within the CESM framework

- Component Integration
- Benchmarking
- Testing
  - Atm. Physics Testing (CESM "F compset")
  - Ocean Testing (CESM "C compset")
- A solid regression baseline will be established on CPUs (Cheyenne) before introducing GPU code.
- Earthworks has received a *Cheyenne* resource allocation to achieve these first-year goals.

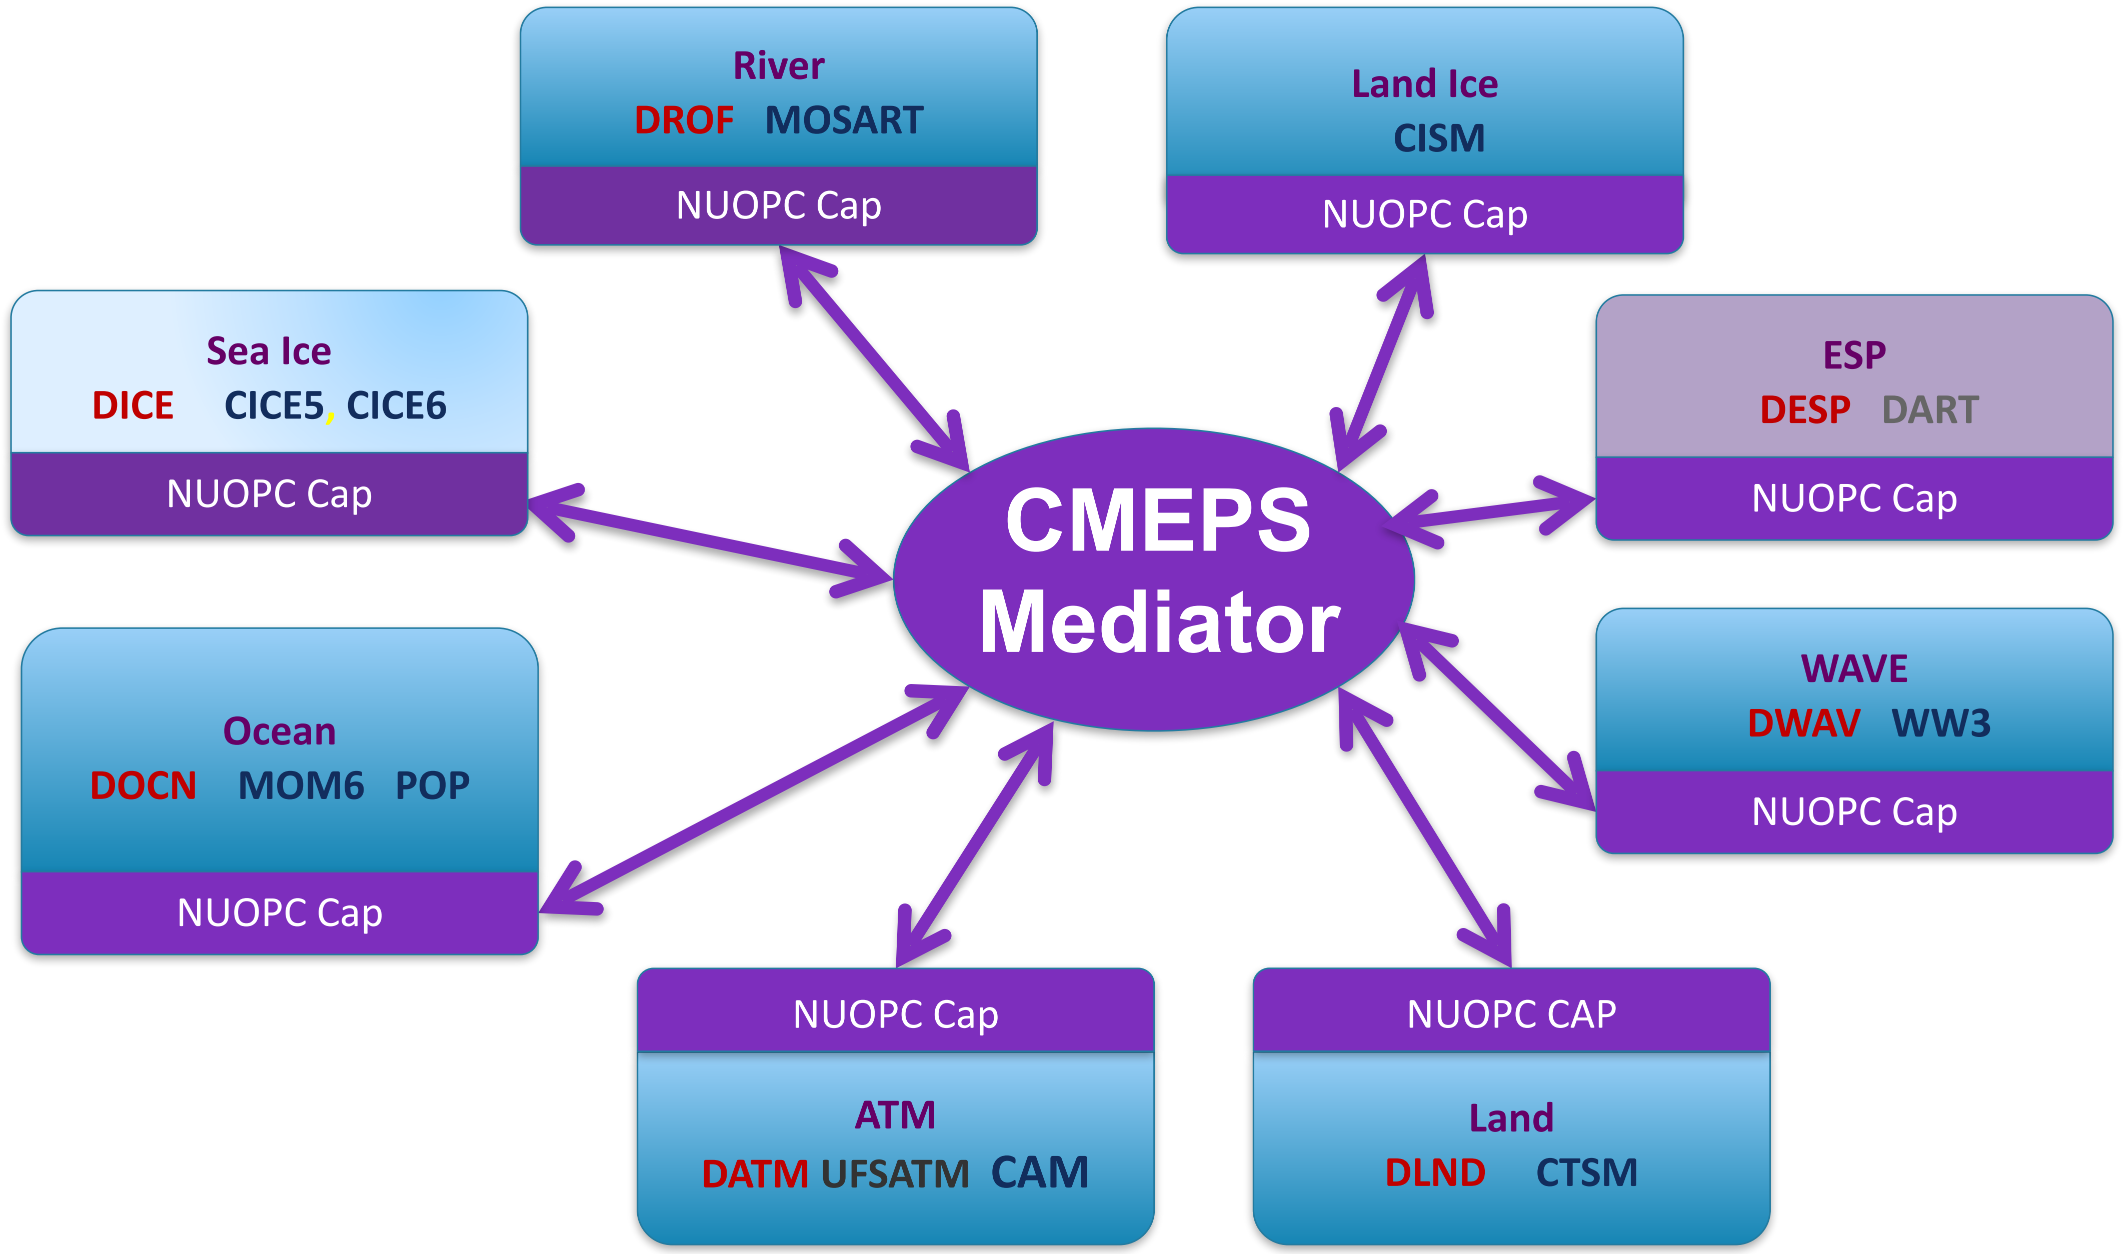

Benefits to CESM as EarthWorks pushes the performance/resolution envelope:

- Increases community experience with exascale technology and accelerates GPU development.
- Accelerates the development of cross-scale atmospheric physics.
- Helps advance incorporating Machine Learning into CESM.
- Even without running the full EarthWorks configuration, university researchers can:
  - *analyze EarthWorks output datasets;*
  - *use EarthWorks software, e.g. big-data tools and the Fortran-callable machine learning interface.*

Education and outreach goals

EarthWorksForce: Partnerships and Student Pipelines

Computational partnerships and collaborations between NCAR, university faculty, students, and vendor experts.

BOISE STATE UNIVERSITY

UNIVERSITY OF DELAWARE · 1743

UNIVERSITY OF WYOMING

Faculty

NCAR NATIONAL CENTER FOR ATMOSPHERIC RESEARCH

NSF

Students

The EarthWorksForce concept: Pipelining students into HPC careers.

- Partner with university professors to integrate performance portability work into curricula.
- Partner with code owner, science stakeholder
- Training materials.
- Leverage near-peer mentoring
- Software Engineers "POD leaders" with both strong HPC and leadership and mentoring skills.

External Advisory Panel

Mariana Vertenstein

Phil Jones

Mark Govett

AJ Lauer

Communication with CESM

EarthWorks will work hard to maintain compatibility with the evolving CESM code base.

We will persistently engage with CESM, and with Community scientists.

In 2025, the future CESM leadership will decide whether or not to adopt EarthWorks as a supported configuration of CESM. We hope that they will adopt it with enthusiasm.

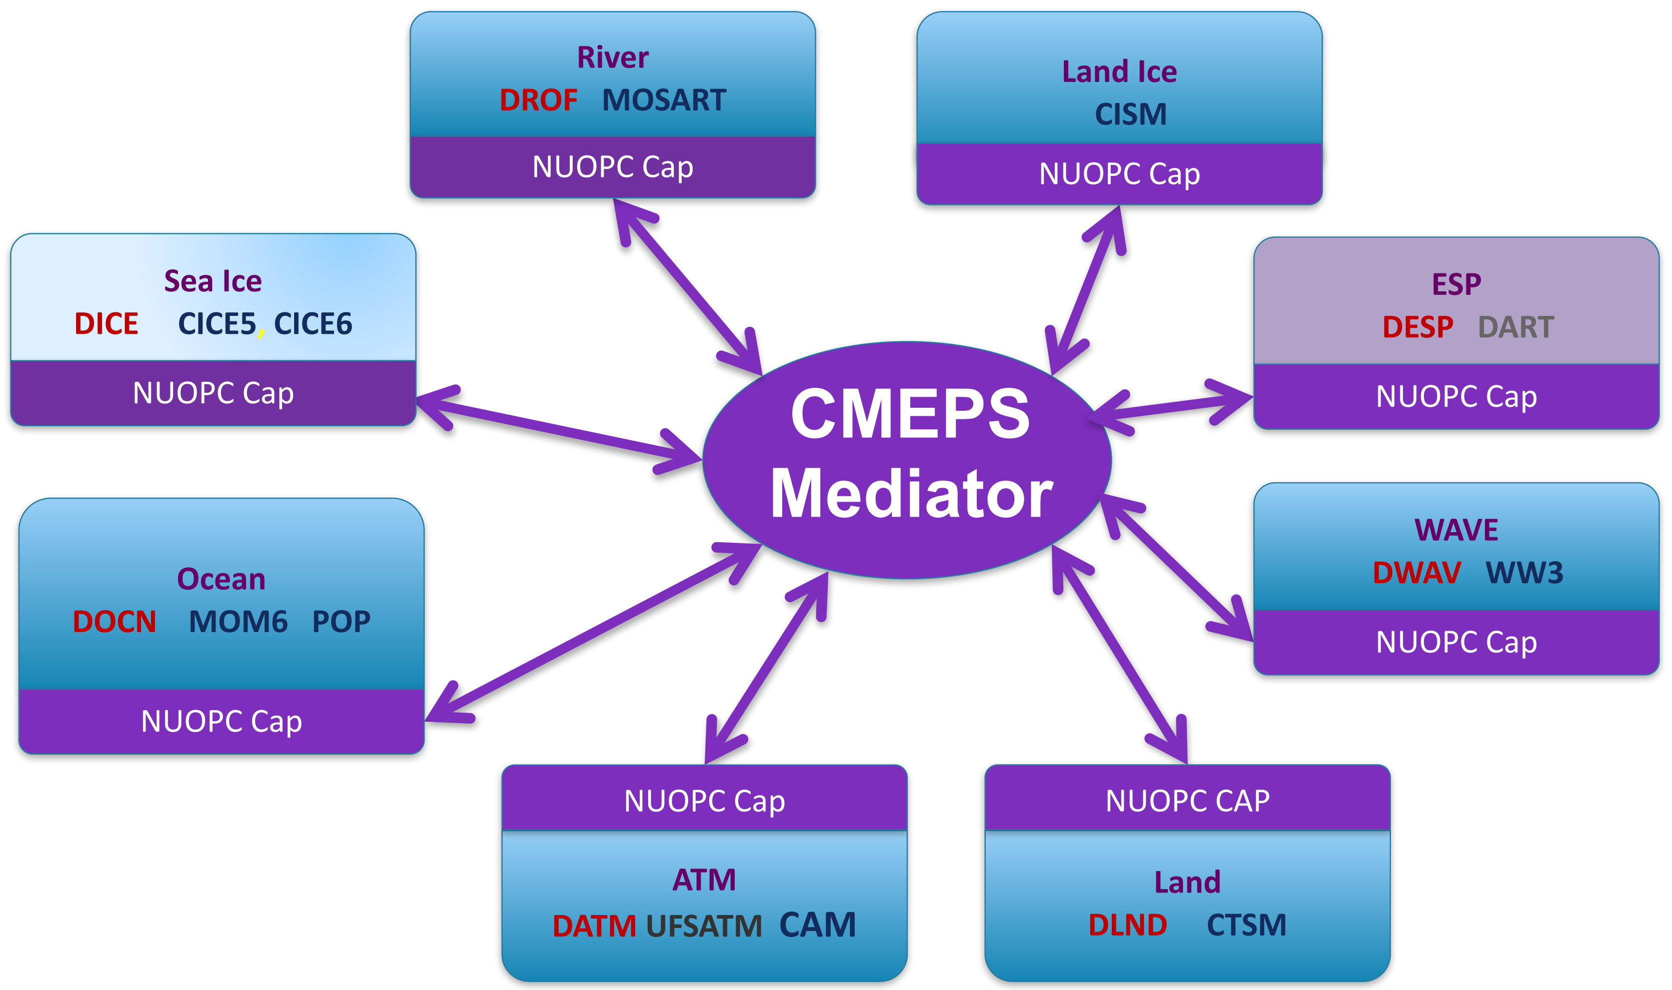

Extra Slides

Workflow

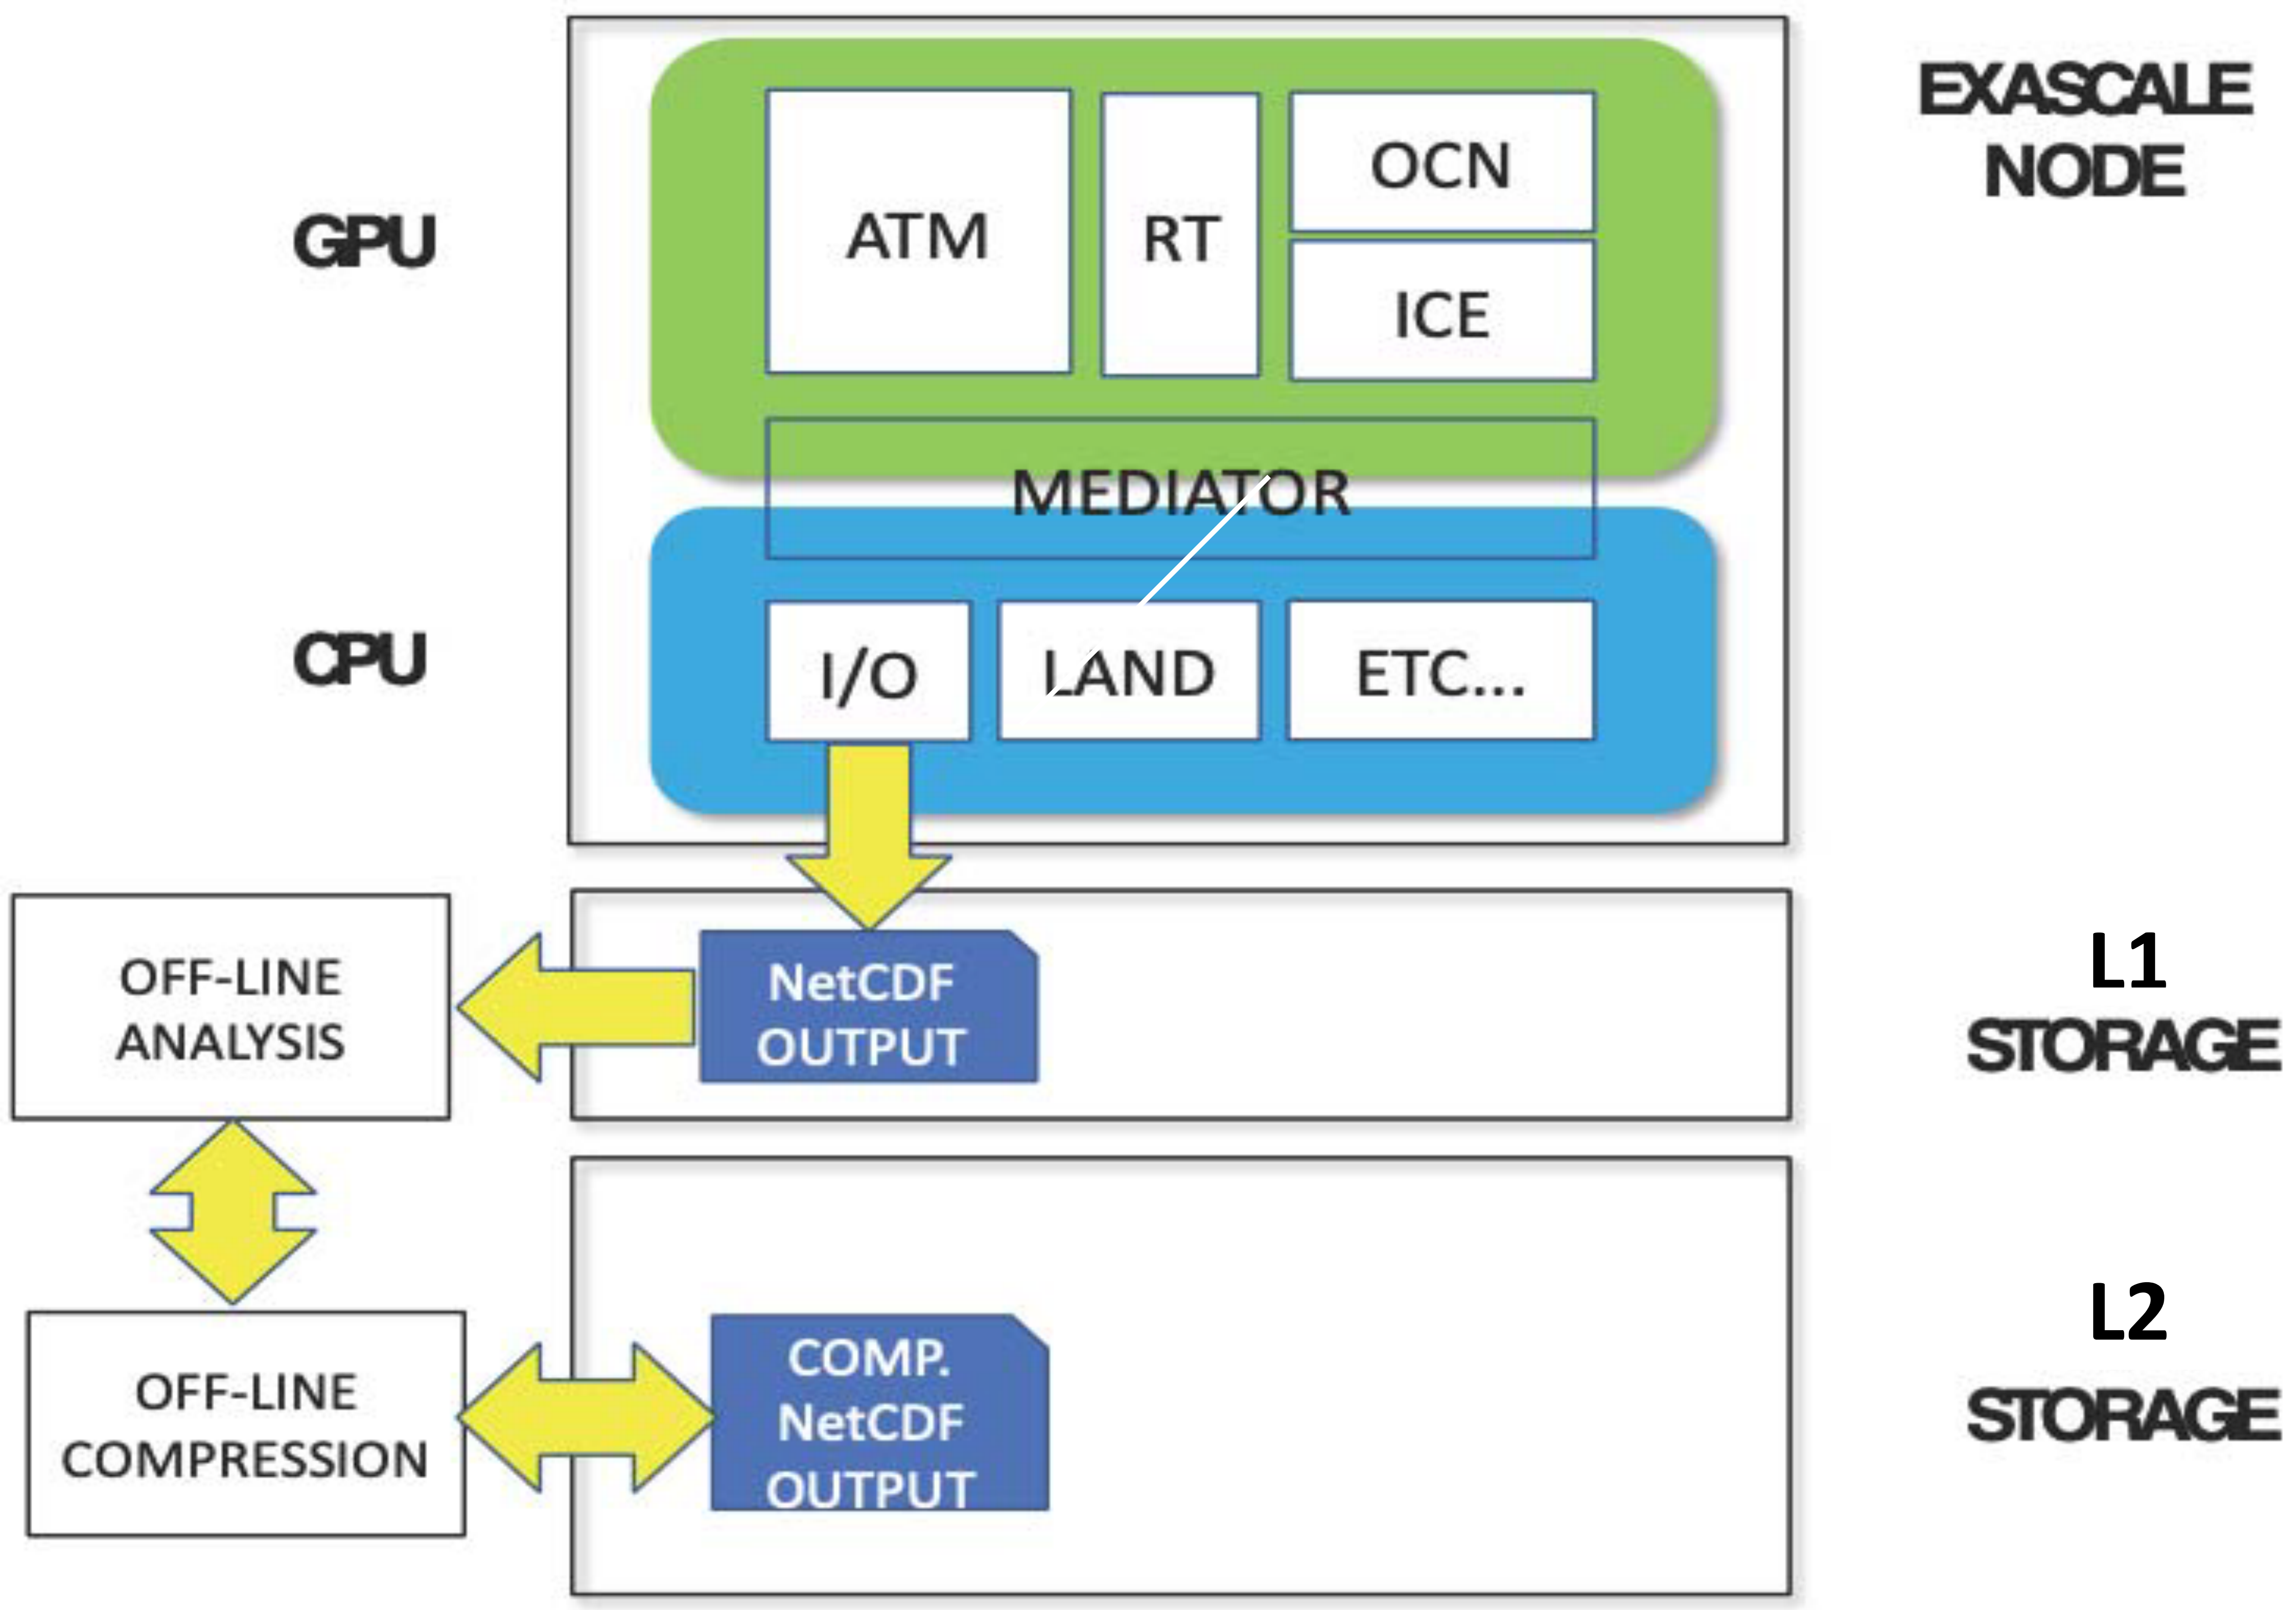

EXASCALE NODE
GPU
ATM
RT
OCN
ICE
MEDIATOR
CPU
I/O
LAND
ETC...
L1 STORAGE
OFF-LINE ANALYSIS
NetCDF OUTPUT
L2 STORAGE
OFF-LINE COMPRESSION
COMP. NetCDF OUTPUT

GPU peak flops and bandwidth are, and have been, outpacing CPUs for some time.

We estimate that GPUs are at least 6x better throughput/Watt on average than CPUs.

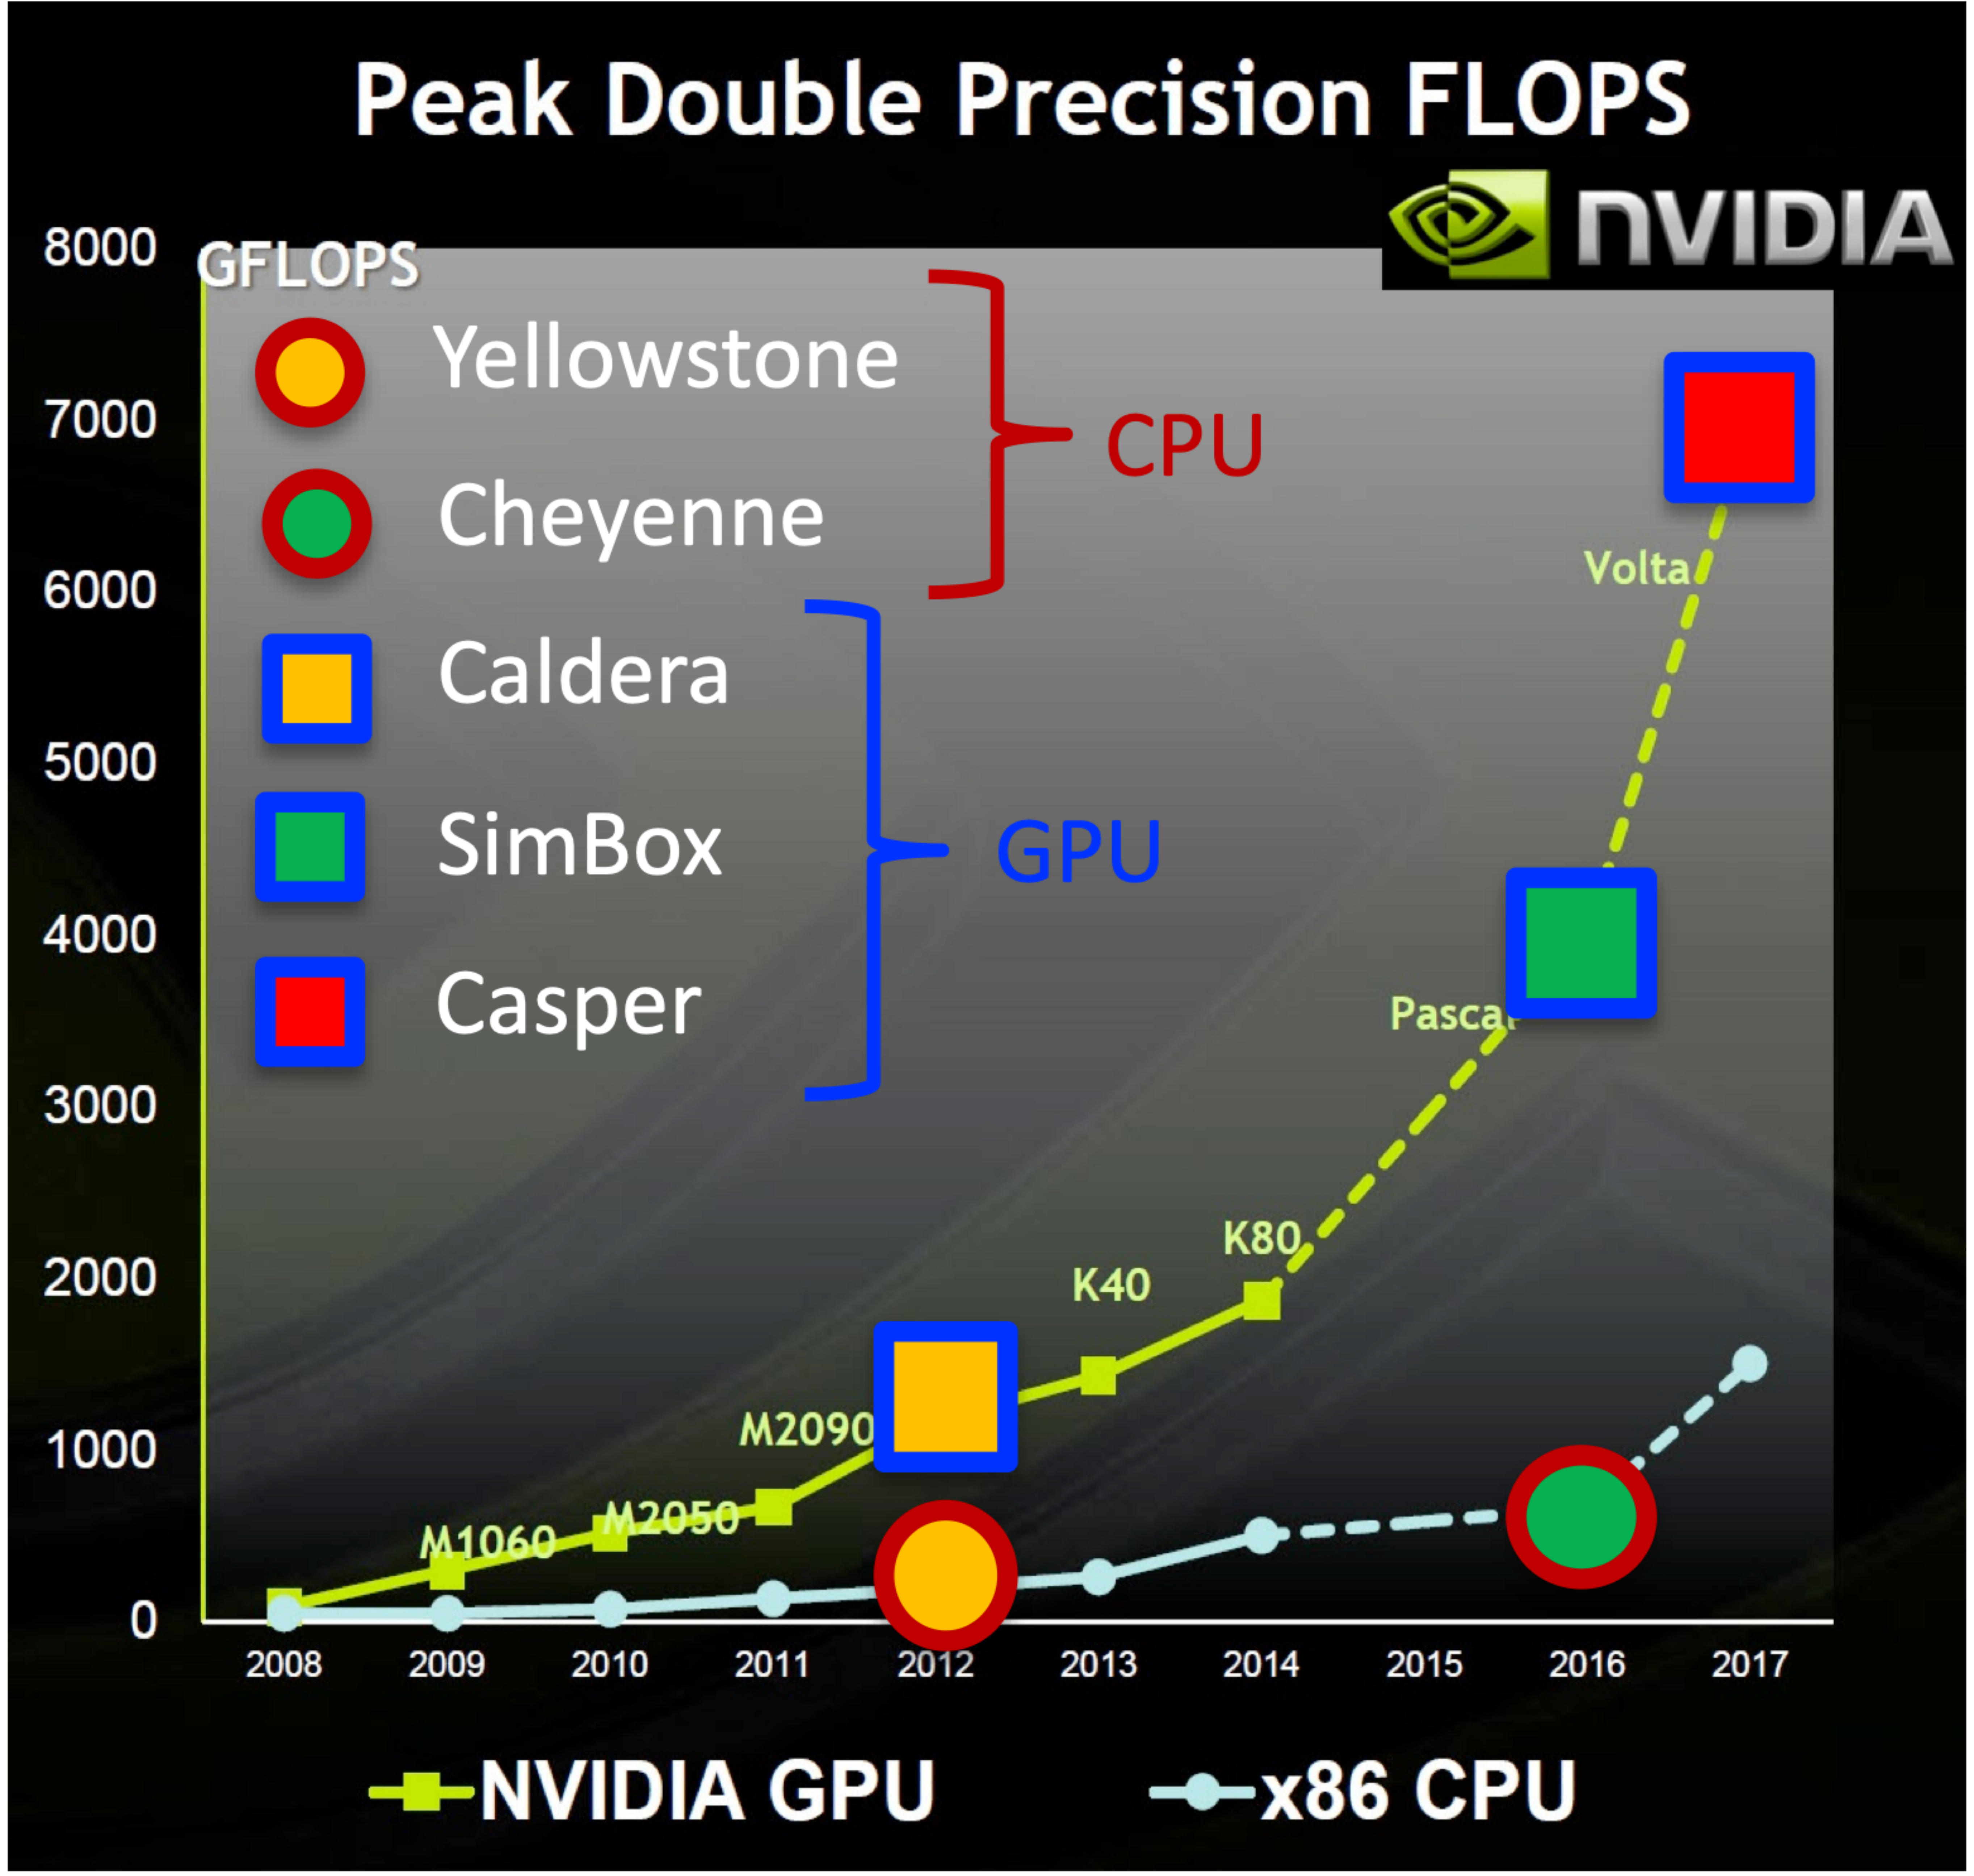

Source: Jensen Huang's (NVIDIA-CEO) IEEE SC15 presentation.

How the pipeline works:

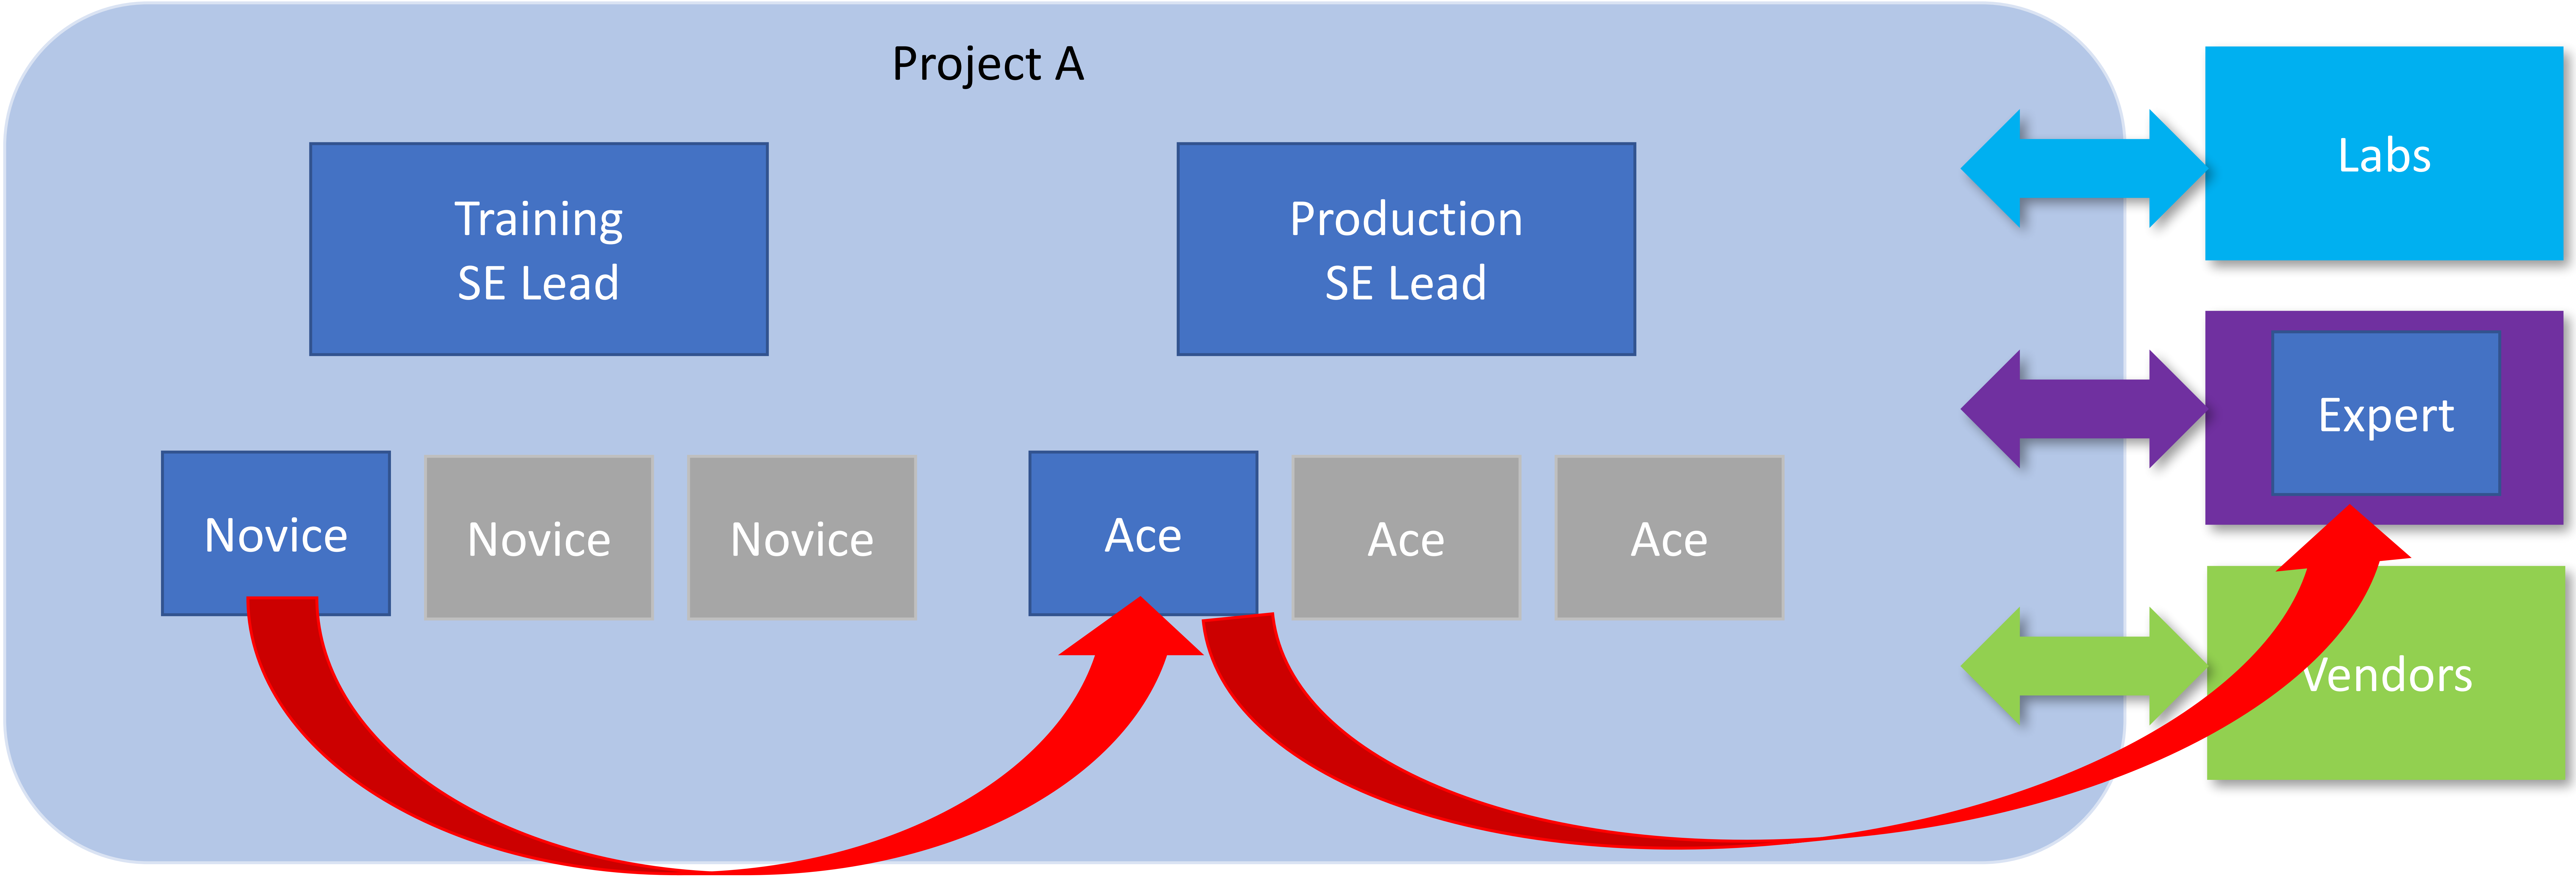

"[Earthworks] relies on training young scientists to be competent in Earth System Model development … Graduating a new batch of good students and post docs who possess this skill could have great impacts on the long-term outlook of climate model development".

Panelist Comment