The High Performance Climate & Weather Benchmark

Xavier Yepes-Arbós & Mario Acosta

(xavier.yepes@bsc.es, mario.acosta@bsc.es)

Barcelona Supercomputing Center (BSC)

17th JLESC workshop, 13 May 2025

EuroHPC Joint Undertaking

Funded by the European Union. This work has received funding from the European High Performance Computing Joint Undertaking (JU) under grant agreement No 101093054.

Co-funded by the European Union

Earth Sciences Department

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

ESiWACE3

Destination Earth

HANAMI

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

Earth Sciences Department | Barcelona Supercomputing Center – Centro Nacional de Supercomputación

ESiWACE3 - Centre of Excellence in Simulation of Weather and Climate in Europe

ESiWACE3 focuses to **support the weather and climate modelling community** to reach **the excellence regarding exascale supercomputing**

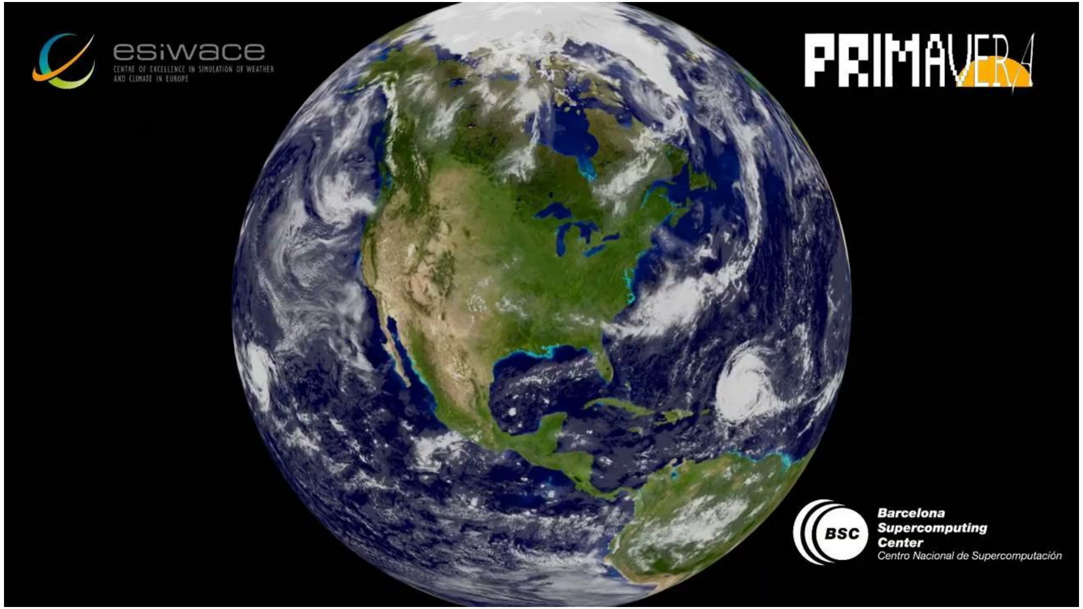

Ocean at different resolutions using the EC-Earth model performed by Oriol Tintó (BSC)

Coordinated Barcelona Supercomputing Center – Centro Nacional de Supercomputación

Consortium of 12 partners from 8 different countries

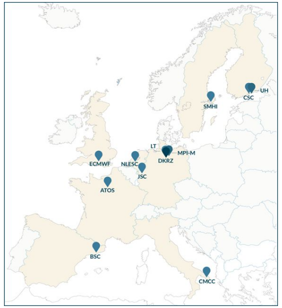

Start: 1 January 2023
End: 31 December 2026

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

ESiWACE3 ecosystem

Models/Tools

ECMWF OpenIFS
NEMO
GridTools 4python
ICON
RISC-V
PSyclone
FVM
AutoRPE
Project ID: 416
HPCW
Field Compression Library
M4M FIP FDP
FDO4Climate
LLVM Compiler Infrastructure
AUTOSUBMIT
Auto Profile
EESSI
Kernel Tuner Ecosystem
Field Compression Benchmark Suite
EC EARTH

HPC Community

MAX
POP
European Processor Initiative (epi)
EUPEX European Pilot for Exascale
THE EUPILOT
HiDALGO Centre of Excellence
HPC3
LUMI
EuroHPC Joint Undertaking
ChEESE
BSC
SPACE
Multiscale
CASTIEL 2
EPICURE Unlocking European-level HPC Support
E-CoE
E4 COMPUTER ENGINEERING
EXCELLERAT
HPC Now!
PerMed CoE
HPC/Exascale Centre of Excellence in Personalised Medicine
EURO CC
ESCAPE 2
bioexcel

ESM Community

Climateurope2
EERIE European Eddy-Rich ESMs
nextGEMS
WARM WORLD
HANAMI
HPC consortium EU - Japan
EuroHPC Joint Undertaking
esm 2025
Destination Earth
EUROPEAN CENTRE FOR MEDIUM-RANGE WEATHER FORECASTS
enes EUROPEAN NETWORK FOR EARTH SYSTEM MODELLING
GANANA
EDITO

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

HPCW - The High Performance Climate & Weather Benchmark

Why a Climate and Weather Benchmark?

HPCW - The High Performance Climate & Weather Benchmark

Why a Climate and Weather Benchmark?

High Performance Linpack

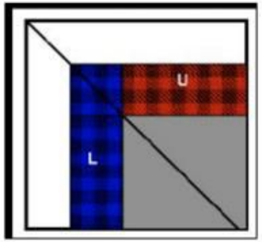

STREAM

STREAM
High Performance Computing

High Performance Conjugate Gradients

HPCG
High Performance Conjugate Gradients

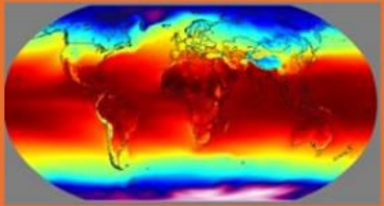

Benchmark adapted for weather & climate considerations ?

1979 | 1996 | 2014

Dense linear equations
Still used for Top500 ranking (and Green500)

Pure memory benchmark

Balanced memory/compute benchmark

Slide provided by Erwan Raffin, Atos

HPCW - The High Performance Climate & Weather Benchmark

Why a Climate and Weather Benchmark?

High Performance Linpack

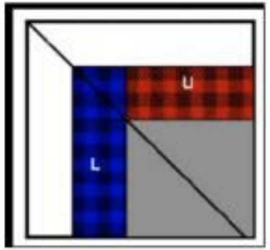

STREAM

STREAM High Performance Computing

High Performance Conjugate Gradients

HPCG High Performance Conjugate Gradients

High Performance Climate & Weather

HPCW

1979 — 1996 — 2014 — 2019 (*2025 OSS*)

- **1979:** Dense linear equations
  Still used for Top500 ranking (and Green500)
- **1996:** Pure memory benchmark
- **2014:** Balanced memory/compute benchmark
- **2019:** Benchmark adapted for weather & climate considerations

Slide provided by Erwan Raffin, Atos

HPCW - The High Performance Climate & Weather Benchmark

What is a benchmark?

Ensuring reproducibility

1. Benchmark
- A code
- A specific version
- A specific configuration
- Its specific dependencies

2. Test-case
- Specific input files (compatible with the code version)
- The reference output files

3. Verification procedure
- Numerical error checking

4. Scoring metrics
- Time to solution
- Gflops
- Energy to solution
- Domain specific metrics (as SYPD, etc.)

Slide provided by Erwan Raffin, Atos

HPCW - The High Performance Climate & Weather Benchmark

What is a benchmark?

Ensuring reproducibility

1. Benchmark
- A code
- A specific version
- A specific configuration
- Its specific dependencies

2. Test-case
- Specific input files (compatible with the code version)
- The reference output files

3. Verification procedure
- Numerical error checking

4. Scoring metrics
- Time to solution
- Gflops
- Energy to solution
- Domain specific metrics (as SYPD, etc.)

How do HPCW's technical choices promote reproducibility?

Slide provided by Erwan Raffin, Atos

Requirements

- Each Weather and Climate model
    - Comes with its build system (Makefile, CMake, . . .)
    - Needs different libraries and tools as dependencies
- Wish list for a framework
    - Build all models and their dependencies the same way
    - Benchmark all models with their relevant test cases
    - Report the results
- But also
    - Simple, easy to use and to maintain
    - Agnostic to
    - each model build system
    - each cluster environment
    - each scheduler system
- Customizable
    - adapt and change dependencies
    - change compilers and flags
    - allow optimizations at all levels

DestinE video example by Oriol Tintó (BSC)

Slide provided by Erwan Raffin, Atos

Agnostic and easy to use for reproducible results

- HPCW is a CMake-based framework
    - able to compile all the “component” on top of their own build system
        - CMake SuperBuild
        - SPACK recipes (optional usage but recommended)
        - CTest
    - agnostic to
        - each code build system (autotools, Makefile, CMake, etc.)
        - each cluster environment (compilers/libraries version, etc.)
        - each scheduler system (slurm, etc.) to launch test cases

Agnostic and easy to use for reproducible results

- HPCW is a CMake-based framework
  - able to compile all the “component” on top of their own build system
    - CMake SuperBuild
    - SPACK recipes (optional usage but recommended)
    - CTest
  - agnostic to
    - each code build system (autotools, Makefile, CMake, etc.)
    - each cluster environment (compilers/libraries version, etc.)
    - each scheduler system (slurm, etc.) to launch test cases

- HPCW can be adapted and customized
  - specificities are managed separately
  - stored in the Git repository as well.

HPCW

Agnostic and easy to use for reproducible results

- HPCW is a CMake-based framework
  - able to compile all the “component” on top of their own build system
    - CMake SuperBuild
    - SPACK recipes (optional usage but recommended)
    - CTest
  - agnostic to
    - each code build system (autotools, Makefile, CMake, etc.)
    - each cluster environment (compilers/libraries version, etc.)
    - each scheduler system (slurm, etc.) to launch test cases

- HPCW can be adapted and customized
  - specificities are managed separately
  - stored in the Git repository as well.

- The advantages of CMake and Spack are
  - to deal with the dependencies
  - and to deal with the dependencies of the dependencies
  - their scripting capabilities for automation

HPCW Roadmap In Europe and beyond

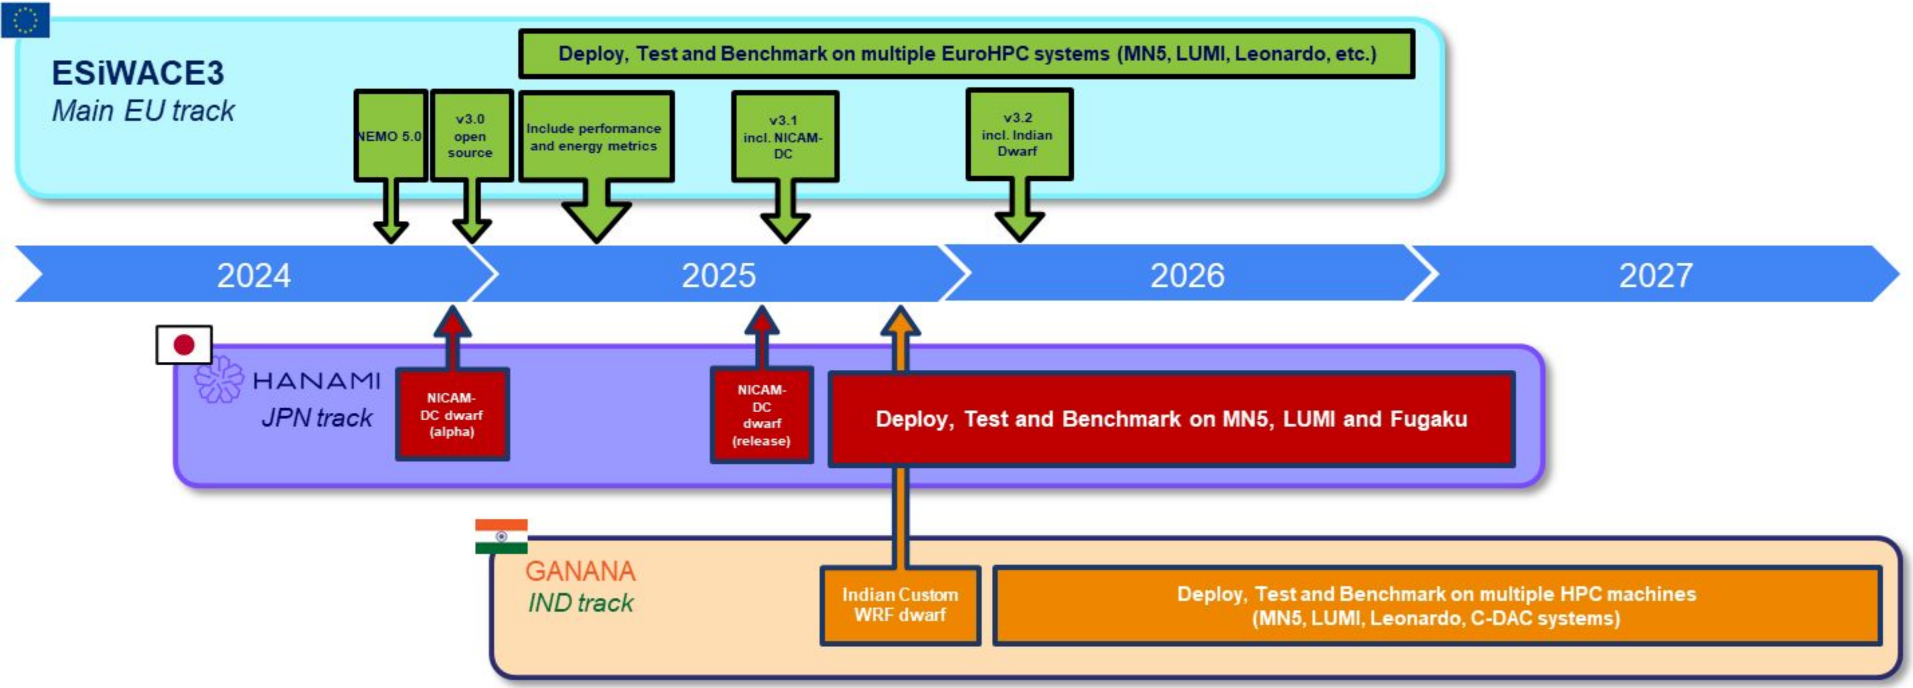

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

AUTOSUBMIT

Autosubmit is a lightweight workflow manager designed to meet **climate research necessities**. Unlike other workflow solutions in the domain, it integrates the capabilities of an **experiment manager**, workflow **orchestrator** and **monitor** in a self-contained application.

Automatization

- Meta-scheduling
- Task dependencies
- High-level config
- Automatic retries

Interoperability

- Multi-platform
- Single point of access
- Python
- Web GUI

Efficiency

- Custom granularity
- Dyn. task aggregation
- Performance metrics

Robustness

- Scalable database
- Manage multiple hosts
- Auto-recovery

Traceability

- Workflow database
- Manage multiple hosts
- Auto-recovery

Monitoring

- Real-time status
- Workflow statistics
- Task logs

AS API

AS GUI

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

D. Manubens-Gil, J. Vegas-Regidor, C. Prodhomme, O. Mula-Valls and F. J. Doblas-Reyes, (2016). "Seamless management of ensemble climate prediction experiments on HPC platforms", 2016 International Conference on High Performance Computing & Simulation (HPCS), Innsbruck, pp. 895-900. https://doi.org/10.1109/HPCSim.2016.7568429

W. Uruchi, M. Castrillo and D. Beltrán, (2021). "Autosubmit GUI: A Javascript-based Graphical User Interface to Monitor Experiments Workflow Execution", Journal of Open Source Software, 6(59), 3049. https://doi.org/10.21105/joss.03049

Web

Cloud

HPC

Workflow integration for Earth-system-model performance assessment

- Develop a workflow (orchestrator and monitor in a self-contained application) for testing three aspects of HPCW and Earth system models (ESMs) in general:
    1. Replicability
    2. Projection skill
    3. Computational performance.
    - Combine those methodologies inside a single workflow with a generic interface applicable to any weather and climate benchmark or model and accessible on any cluster.

Earth System Model (ESM) performance assessment

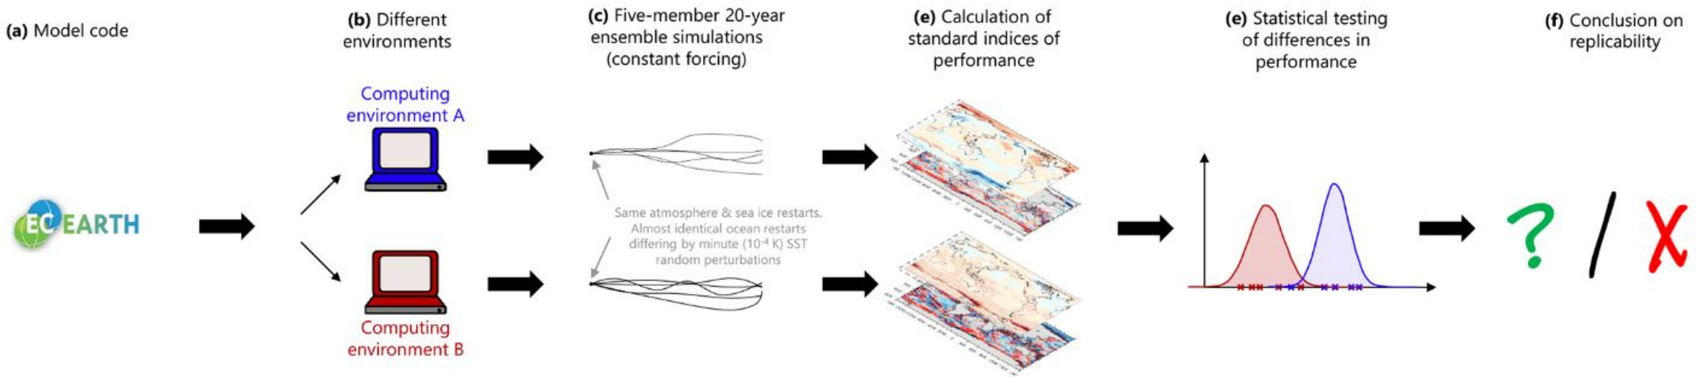

Massonnet, F., Ménégoz, M., Acosta, M., Yepes-Arbós, X., Exarchou, E., & Doblas-Reyes, F. J. (2020). **Replicability of the EC-Earth3 Earth system model under a change in computing environment.** Geoscientific Model Development, 13(3), 1165–1178. https://doi.org/10.5194/gmd-13-1165-2020

- ***Testing three aspects of ESMs***
  - ***I. Replicability***
  - ***II. Projection skill***
  - ***III. Computational performance***

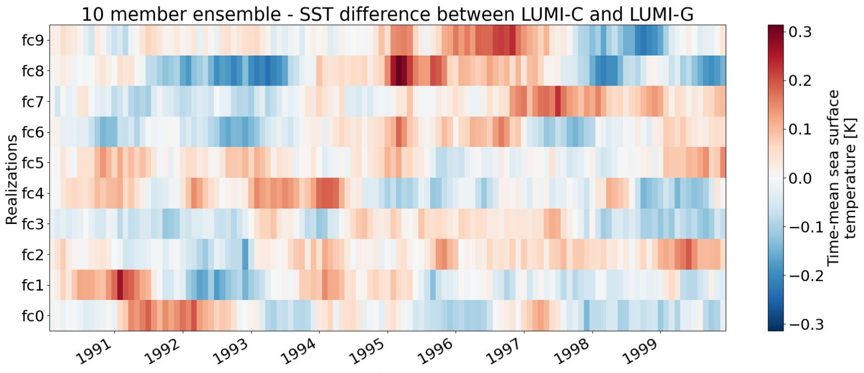

Earth System Model (ESM) performance assessment

AWI-CM3 CMPI: 0.931

	arctic MAM	arctic JJA	arctic SON	arctic DJF	northmid MAM	northmid JJA	northmid SON	northmid DJF	tropics MAM	tropics JJA	tropics SON	tropics DJF
siconc	1.07	0.92	0.74	0.98	0.67	1.07	1.23	0.84				
tas	1.08	0.42	0.42	0.80	0.63	0.97	0.92	0.69	0.79	0.62	0.70	0.75
clt	0.90	1.16	1.19	1.07	0.70	0.76	0.66	0.78	0.85	0.78	0.61	0.68
pr	0.77	0.87	1.02	1.07	0.87	1.22	1.10	0.91	1.11	1.00	0.90	0.84
rlut	1.02	0.88	0.61	0.49	1.01	0.67	0.63	0.90	1.21	1.04	0.95	0.92
uas	0.65	0.85	0.64	0.91	0.70	0.84	0.56	0.67	0.98	0.80	0.81	0.70
vas	0.62	0.80	0.69	0.82	0.73	0.74	0.74	0.65	0.95	0.80	0.81	0.74
300hPa ua	0.75	0.99	0.95	1.23	0.94	1.23	0.86	1.12	0.95	0.73	0.77	0.82
500hPa zg	0.29	0.62	1.31	1.05	0.41	0.85	0.78	0.63	0.38	0.27	0.61	0.31
SD zos	0.66	0.42	0.61	0.61	0.88	0.93	0.90	0.87	1.02	1.02	1.06	1.07
SD tos	1.06	1.04	0.97	1.06	0.81	1.77	1.49	0.97	1.07	1.10	0.89	1.20
mlotst	1.25	0.55	0.70	1.09	2.32	0.64	0.86	1.85	1.41	1.02	1.54	1.01

	nino34 MAM	nino34 JJA	nino34 SON	nino34 DJF	southmid MAM	southmid JJA	southmid SON	southmid DJF	antarctic MAM	antarctic JJA	antarctic SON	antarctic DJF
siconc					0.97	1.15	1.23	1.16	1.79	3.49	3.50	1.68
tas	0.35	1.06	0.70	0.30	1.64	1.35	1.44	1.60	1.66	1.59	1.23	0.80
clt	0.91	0.32	0.56	0.74	0.80	0.99	0.87	0.79	0.97	0.97	0.83	0.82
pr	1.38	0.91	0.96	1.00	1.17	0.73	1.10	1.08	1.11	1.19	1.04	0.79
rlut	1.52	0.86	0.80	1.04	0.54	0.53	0.59	0.63	1.30	1.44	0.69	0.80
uas	1.34	1.12	0.40	0.70	0.79	0.67	0.78	0.41	0.50	0.45	0.40	0.41
vas	1.21	1.25	0.81	0.65	0.88	0.93	0.68	0.47	0.52	0.50	0.39	0.44
300hPa ua	0.60	0.85	0.38	0.30	0.67	0.91	1.02	0.46	0.91	0.85	0.75	0.76
500hPa zg	0.26	0.22	0.67	0.31	0.94	0.67	0.46	0.74	0.62	0.34	0.18	0.77
SD zos	1.35	1.44	1.71	1.70	1.02	1.05	1.08	1.04	0.95	0.82	0.97	0.97
SD tos	0.28	0.19	0.42	0.52	1.06	0.96	1.01	1.12	0.84	1.73	1.67	0.77
mlotst	0.61	0.66	0.89	0.62	0.51	1.89	2.50	1.08	1.96	2.85	2.81	3.64

Testing three aspects of ESMs

I. Replicability
*II. **Projection skill***
III. Computational performance

Reichler & Kim indices
a.k.a.
Performance scores

Streffing, J., Sidorenko, D., Semmler, T., Zampieri, L., Scholz, P., Andrés-Martínez, M., Koldunov, N., Rackow, T., Kjellsson, J., Goessling, H., Athanase, M., Wang, Q., Hegewald, J., Sein, D. V., Mu, L., Fladrich, U., Barbi, D., Gierz, P., Danilov, S., … Jung, T. (2022). **AWI-CM3 coupled climate model: Description and evaluation experiments for a prototype post-CMIP6 model.** Geoscientific Model Development, 15(16), 6399–6427. https://doi.org/10.5194/gmd-15-6399-2022

Earth System Model (ESM) performance assessment

Metric	Description
Simulation Year Per Day (SYPD)	Simulated years per day in a 24 h period, collected by timing a segment of a production run of usually one year.
Core-hours Per year Simulated (CHSY)	Simulated years produced with respect to the number of parallel resources used
Complexity (Cmpx)	number of prognostic variables per component
Actual SYPD (ASYPD)	how queue time and interruptions affect the complete experiment duration
Parallelization (Paral)	total number of cores allocated for the run
Energy Cost Per Year (JPSY)	Energy in Joules needed per year of simulation
Memory Bloat (Mem B)	ratio between actual and ideal memory size
Data Output Cost (DO)	time and resources used for performing I/O. The value is given as the percentage added to the simulation without outputs. For example, 1.05 means that DO is 5%.
Data Intensity (DI)	amount of data produced in GB per compute-hour
Coupling Cost (Coup C)	time and resources used in the execution of the coupling algorithm as well as load imbalance among model components. The value is given as the percentage represented comparing to the simulation of the components without coupling. For example, 0.05 means that Coup. C. is 5%.

- ***Testing three aspects of ESMs***
  - *I. Replicability*
  - *II. Projection skill*
  - ***III. Computational performance***

Mario Acosta et al.2024. The computational and energy cost of simulation and storage for climate science: lessons from CMIP6. Geoscientific Model Development (GMD). https://doi.org/10.5194/gmd-2023-188

esiwace

CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

Thank you!

EXCELENCIA SEVERO OCHOA

MINISTERIO DE CIENCIA, INNOVACIÓN Y UNIVERSIDADES

AGENCIA ESTATAL DE INVESTIGACIÓN

This work was supported by the Severo Ochoa Program (grant CEX2021-001148-S) funded by MICIU/AEI/10.13039/501100011033

17th JLESC workshop, 13 May 2025

EuroHPC Joint Undertaking

Funded by the European Union. This work has received funding from the European High Performance Computing Joint Undertaking (JU) under grant agreement No 101093054.

Co-funded by the European Union

BSC Barcelona Supercomputing Center
Centro Nacional de Supercomputación