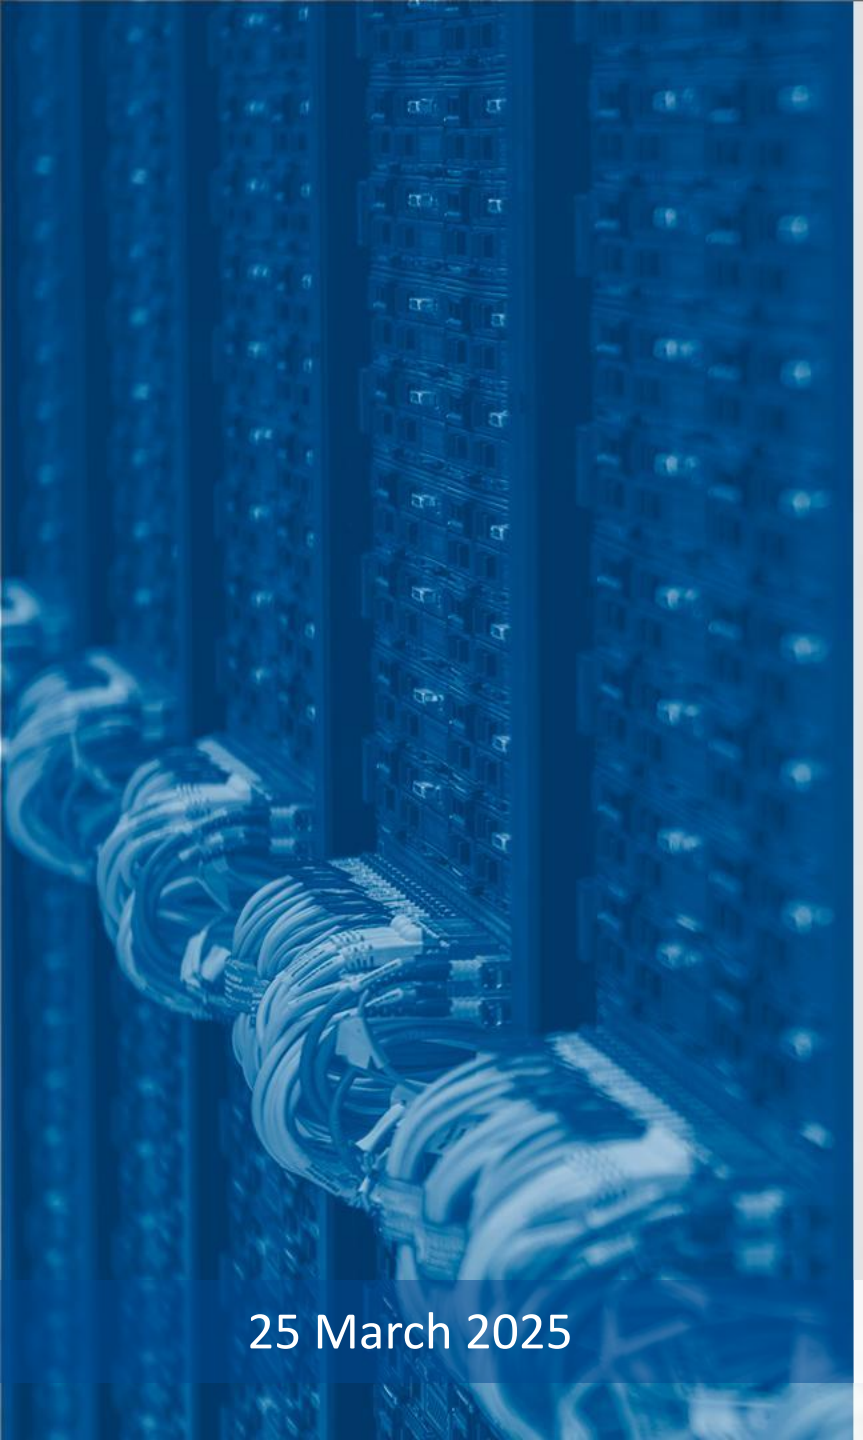

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

HPC adaptation and developments for EC-Earth4 at BSC

Xavier Yepes-Arbós, Mario Acosta, Marta Alerany, Marcel Cases, Maitane Fariñas, Kai Keller, Sergi Palomas, Roc Salvador, Oriol Tintó

25 March 2025

EC-Earth GA 2025, Stockholm, Sweden

HPC for Earth Sciences Team at BSC

- Knowledge about:
  - Mathematical and computational aspects of Earth system models (ESMs)
  - HPC adaptation and development
- Support on computational performance analysis and optimization
- Research on different HPC topics, including GPUs, reproducibility, RISC-V co-design, etc.

CPMIP

- New **CMIP Task team**: Collection CPMIPs and Carbon Footprint for CMIP7

Tier	Level	Description	Metrics
01	Institution	Overall consumption and platform information	• Number of simulated years • Core-hours consumed • Data output • Energy cost (Carbon footprint) • HPC platform
02	Experiment	Metrics by experiment that require close to no effort to obtain	• Simulated years per day (SYPD) • Actual SYPD • Core-hours per simulated year (CHSY) • Joules per simulated year • Parallelisation • Data Intensity • Resolution
03	Experiment	Detailed metrics, not required but collected from institutions willing to collaborate (e.g. ENES)	• Data output cost • Coupling Cost • Memory Bloat • Complexity

CPMIP

- Outcome → Multi-model multi-platform performance study
  - **Supervise & standardise** the collection for CMIP7
  - **Analyse & publish** results (ESGF?)
  - **Improve performance** by identifying machine/model bottlenecks)

- EC-Earth4
  - Be ready for the collection
  - **Automatisation** is key (e.g. auto-ece4)
  - Collect during spin-up runs -> **increased throughput** production

PERFORMANCE METRICS

WARNINGS (2)

SUMMARY

Parallelization: 264 **RSYPD**: 1.56

Metric	Value	Min	Max	SD	MAD
JPSY	140,176,222.22	44,887,700.00	270,110,950.00	89,349,513.24	80,734,485.19
SYPD	0.99	0.49	3.53	1.18	1.06
ASYPD	0.75	0.43	2.82	0.77	0.59
CHSY	6,403.18	1,793.37	13,035.37	4,716.60	4,320.34

Value: Value of the metric calculated at the experiment level.
SD: Standard Deviation.
MAD: Mean Absolute Deviation Around the Mean.

CONSIDERED JOBS

 # considered: 9 **# not considered**: 0

Chunk	Job Name	Queue	Run	CHSY	SYPD	ASYPD	JPSY	Energy
1	a8o8_20191228_000_1_SIM	00:00:24	00:08:06	13,008.60	0.49	0.46	257,135,200	704,480
1	a8o8_20191228_001_1_SIM	00:00:24	00:08:07	13,035.37	0.49	0.46	256,587,700	702,980
1	a8o8_20191228_002_1_SIM	00:00:43	00:03:14	5,192.73	1.22	1.00	117,402,250	321,650
2	a8o8_20191228_000_2_SIM	00:01:39	00:01:59	3,185.23	1.99	1.09	81,975,350	224,590
2	a8o8_20191228_001_2_SIM	00:01:39	00:02:57	4,737.70	1.34	0.86	123,026,900	337,060
2	a8o8_20191228_002_2_SIM	00:01:17	00:07:51	12,607.10	0.50	0.43	270,110,950	740,030
3	a8o8_20191228_000_3_SIM	00:00:49	00:01:07	1,793.37	3.53	2.04	44,887,700	122,980
3	a8o8_20191228_001_3_SIM	00:04:26	00:01:24	2,248.40	2.82	0.68	58,710,250	160,850
3	a8o8_20191228_002_3_SIM	00:00:16	00:01:08	1,820.13	3.48	2.82	51,749,700	141,780

Automatic report by Autosubmit

Load-balance for Coupled ESMs

Load-balance is one key limiting factors of coupled ESMs performance.
Average CMIP6 12.8% (Acosta et. al. 2024)

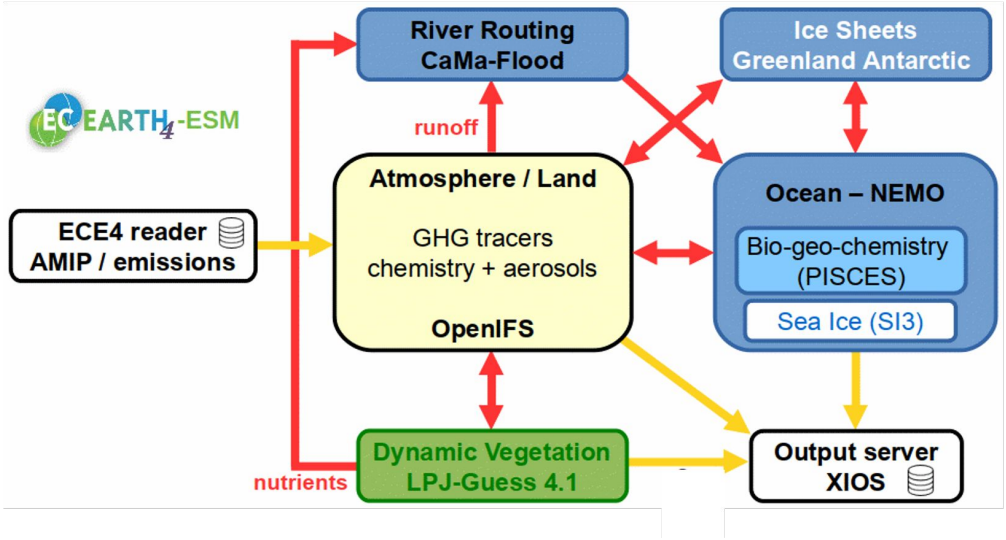

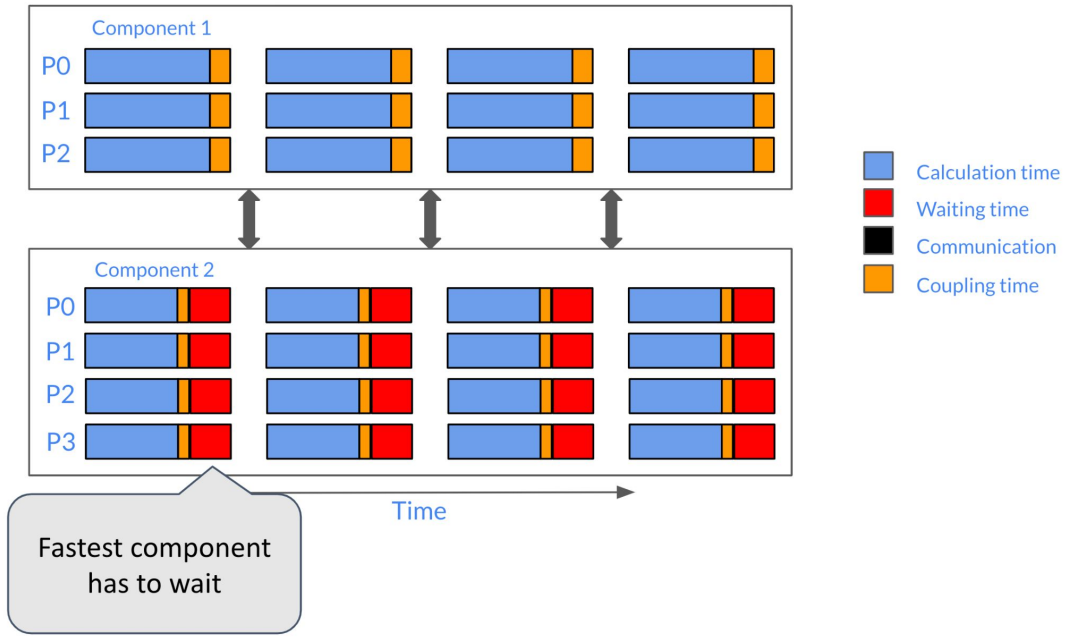

Coupling → Synchronization → Load imbalance → **Parallel inefficiency**

Load-balance for Coupled ESMs

Auto Load-Balance Workflow

- A tool to automatically **minimize the performance loss on the coupling** for ESMs by:
  - Minimizing synchronization cost
  - Finding a **good resource configuration** for the components based on specified Energy-to-Solution / Time-to-Solution criteria
- Easy to adopt by the climate community → integrated into climate workflows, directly compatible with OASIS Coupler
- Tested **portability** to HPC platforms (HPC2020, MeluXina, MareNostrum 4, MareNostrum 5)

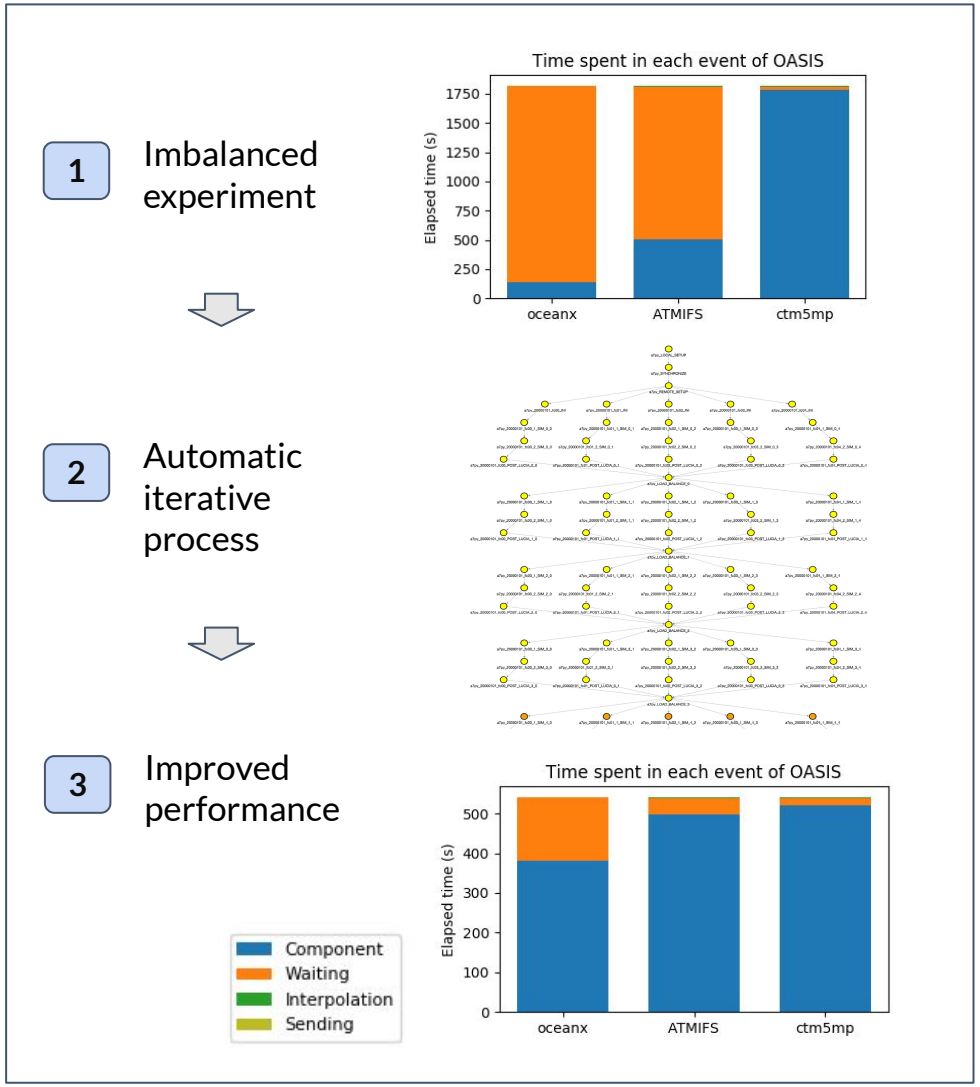

Profiling and optimization

Preliminary scalability of EC-Earth4 - OpenIFS 48 on MareNostrum 5

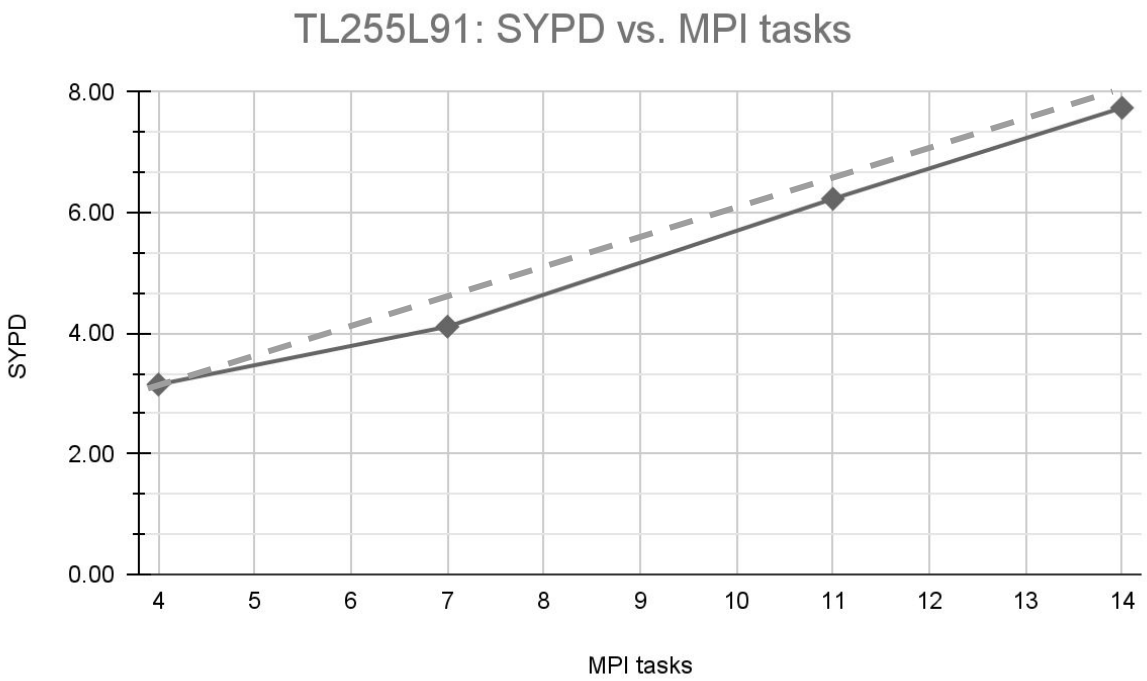

One month simulation of OpenIFS 48 (TL255L91)
In 1 node, 8 OpenMP threads per task

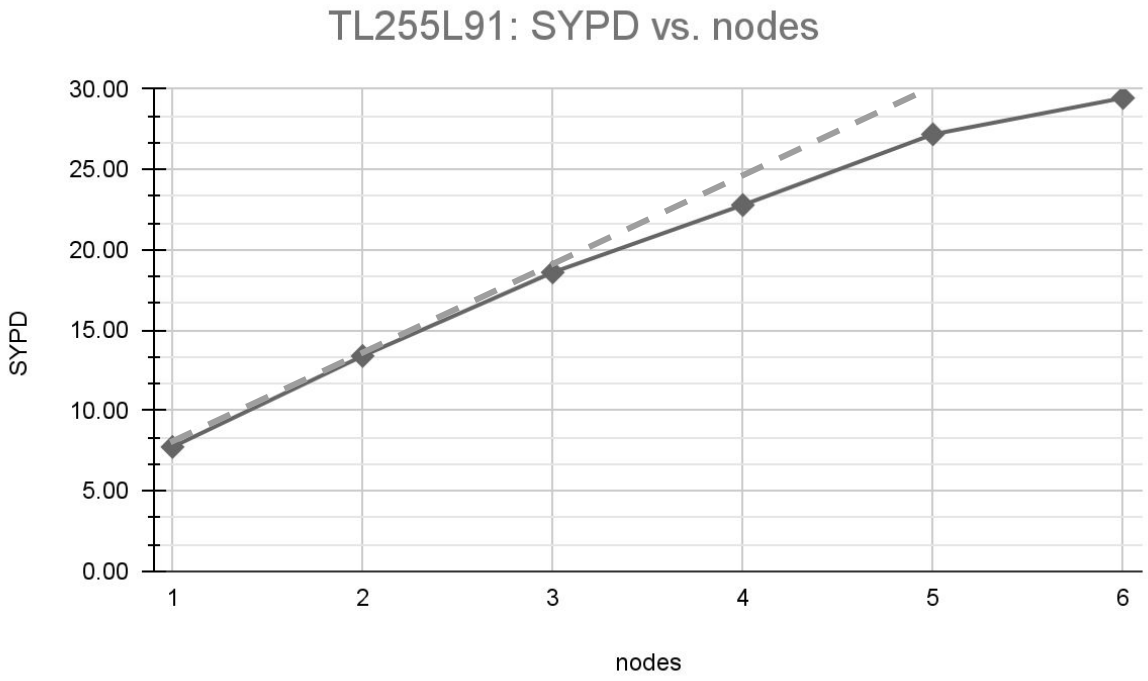

One month simulation of OpenIFS 48 (TL255L91)
14 MPI tasks per node, 8 OpenMP threads per task

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Profiling and optimization

Preliminary profiling of EC-Earth 4 on MN5 running TL255L91 and eORCA1.

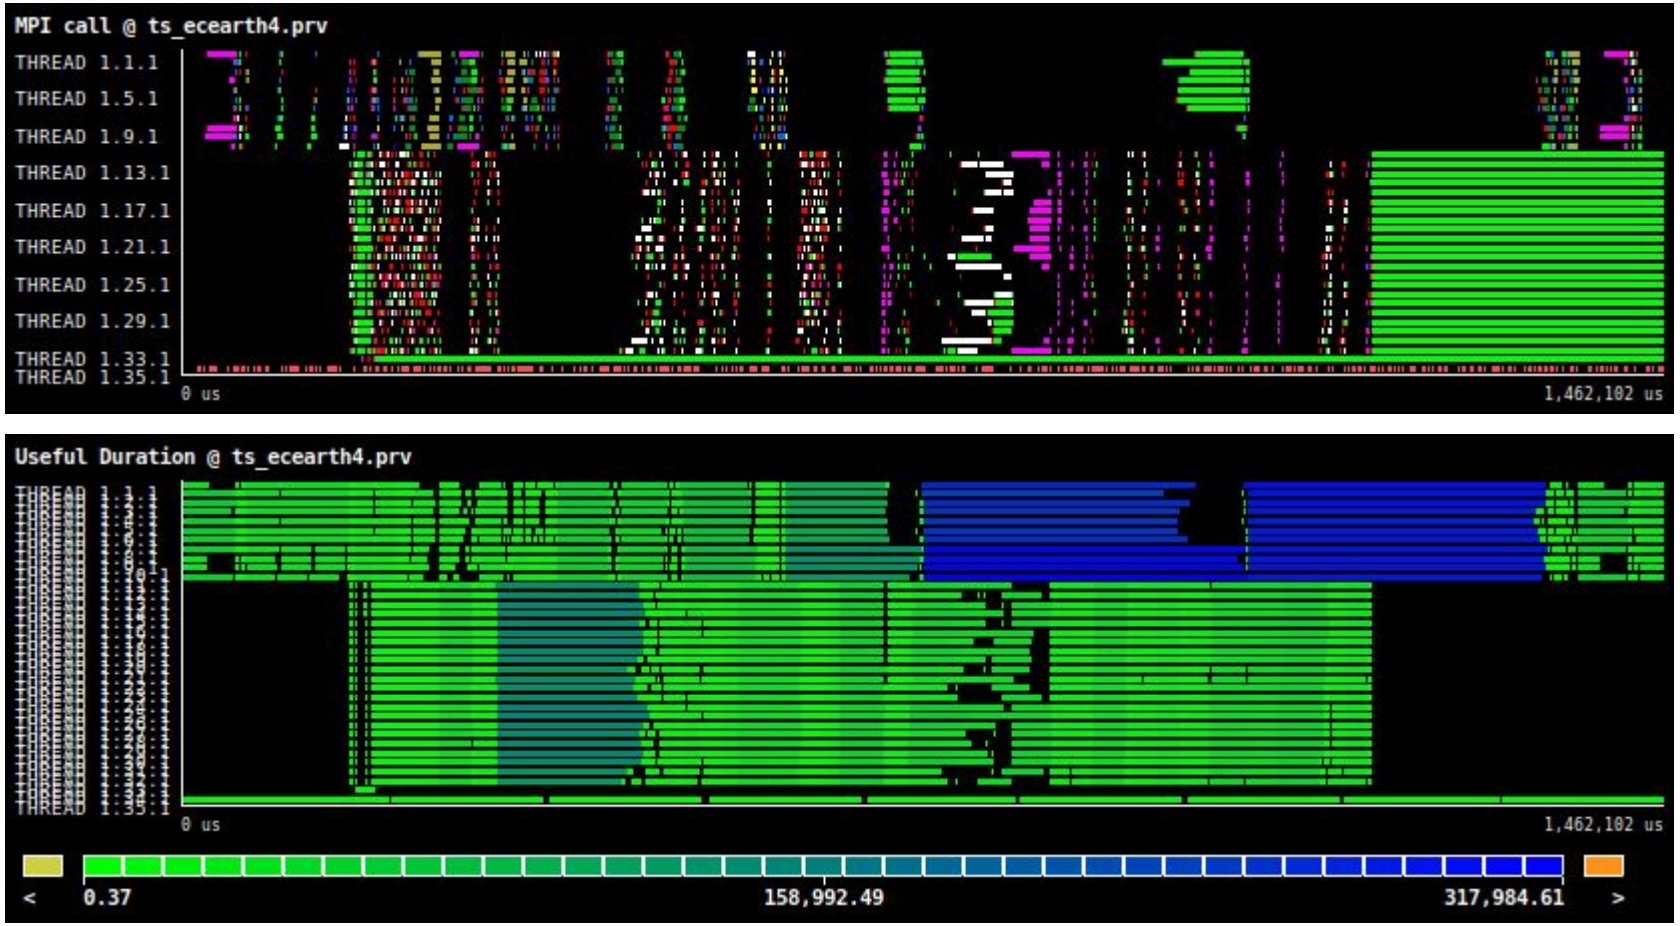

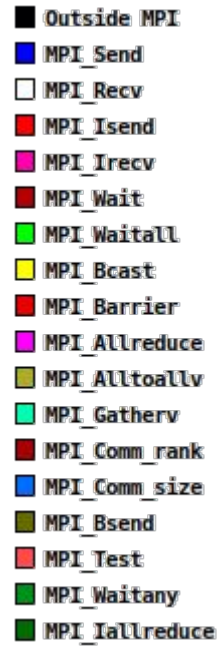

M7 mixed precision

Objective:

To determine which parts of the code can effectively use lower precision without compromising the quality of the results.

Followed steps and challenges:

- **OpenIFS - M7:**
  - Compilation: building the code inside a container and on MareNostrum5.
  - Running the tests: gathering the necessary data for the tests and successfully running the tests in both environments.
- **AutoRPE tool:**
  - Preprocessed sources: AutoRPE works on code that can be compiled.
    - ⋆ *Different build-system makes harder to obtain the preprocessed files*
  - Adapting the tool to new code:
    - ⋆ *Identify the files to be modified by the tool*
    - ⋆ *Adapt the tool to different Fortran coding styles*

Climate Model Replicability

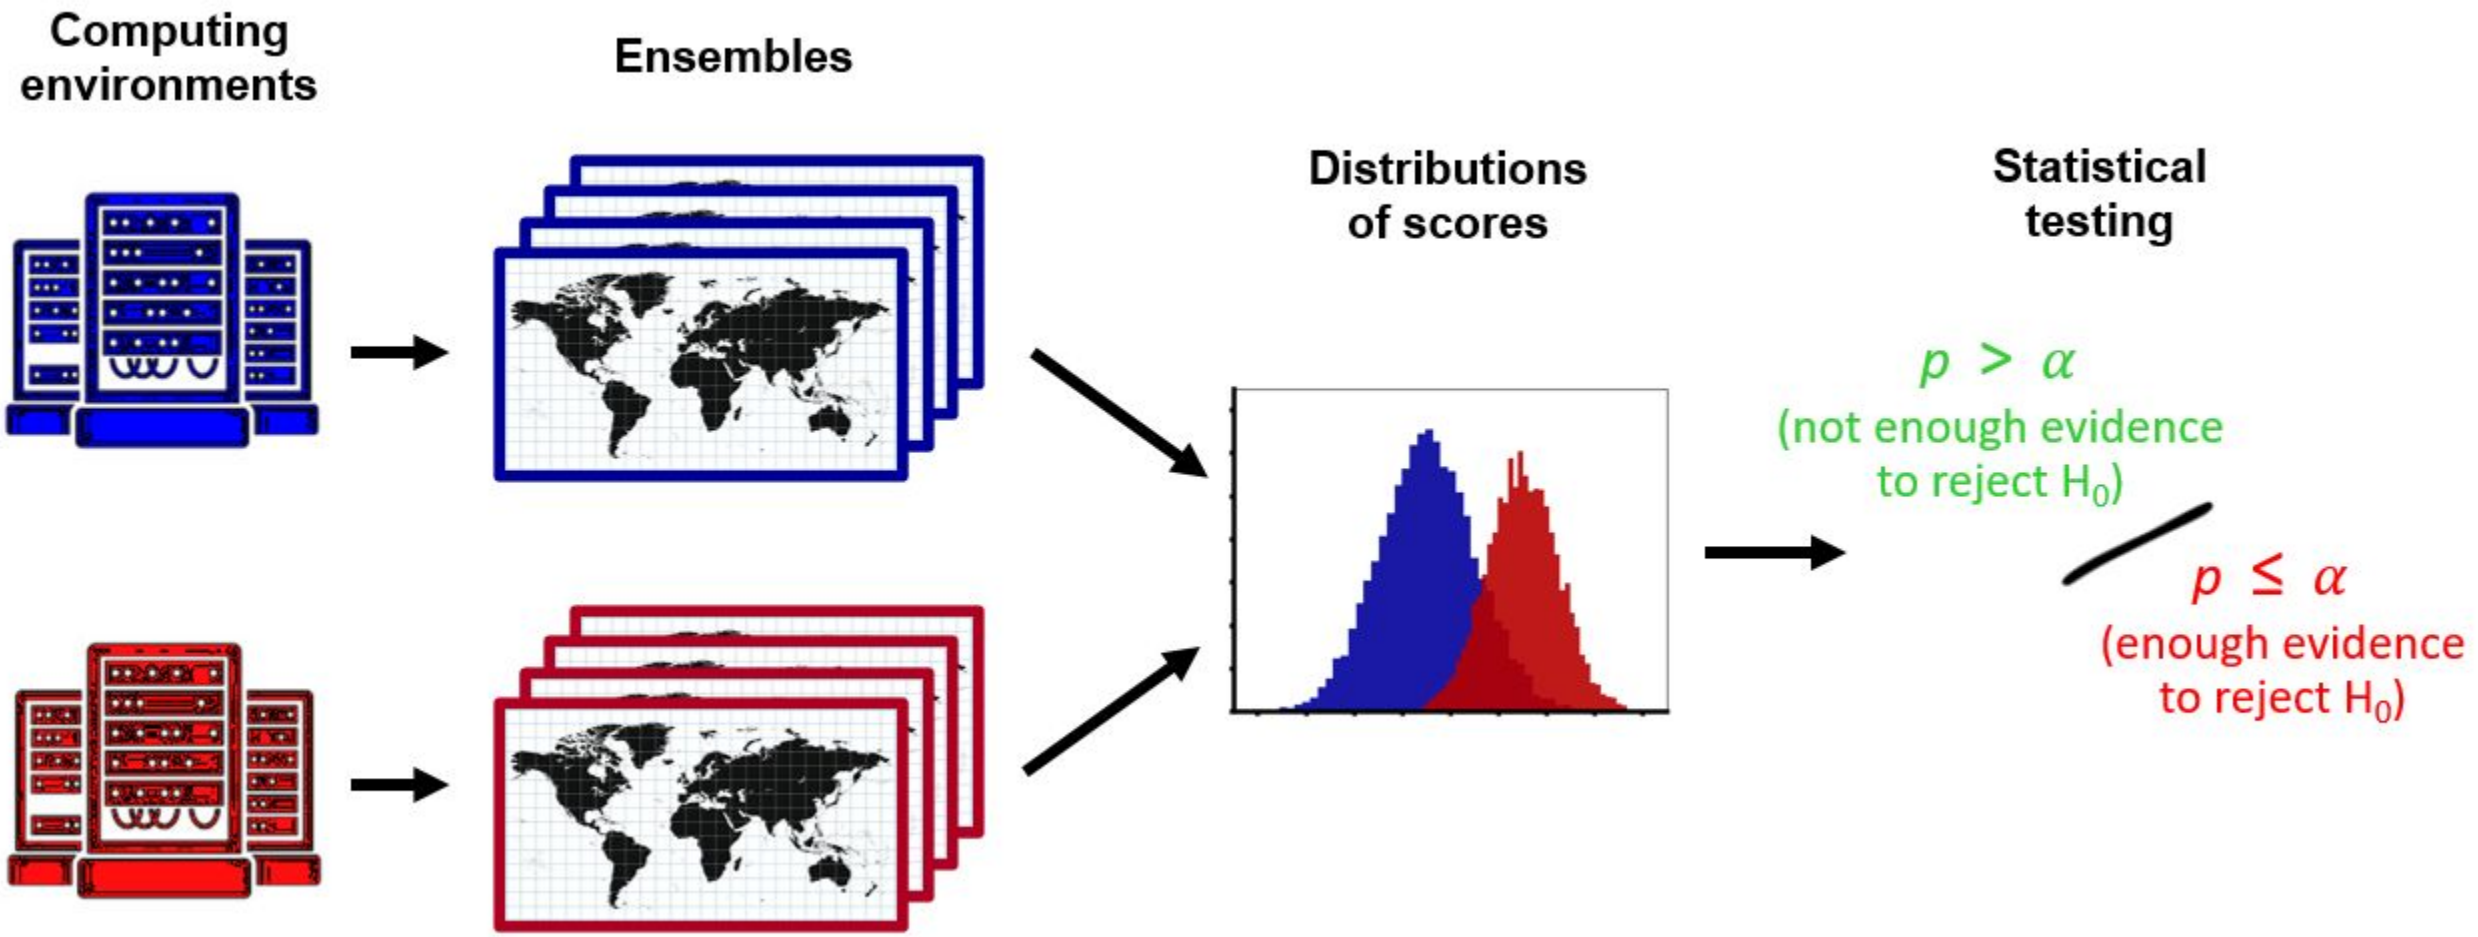

Climate Model Replicability

Improved methodology, based on Massonnet et al., 2022 (https://doi.org/10.5194/gmd-13-1165-2020)

Metrics

- **RK08 score**

$$eRK08_m^X = \sqrt{\sum_{i=1}^{N} \widetilde{\omega}_i \cdot \exp\left[-\left(\frac{x_{mi} - \bar{y}_i}{\sigma_i^y}\right)^2\right]}.$$

- **Bias score**

$$eBIAS_m^X = \sum_{i=1}^{N} \widetilde{\omega}_i \cdot e^{-|\mathrm{bias}_{mi}^X|/\sigma_i^y}.$$

- **RMSE score**

$$eRMSE_m^X = \sum_{i=1}^{N} \widetilde{\omega}_i \cdot e^{-\mathrm{crmse}_{mi}^X/\sigma_i^y}$$

- **combined score**

$$\mathrm{COMBINED} = \frac{eRK08_m^X + eBIAS_m^X + eRMSE_m^X}{3}$$

Tests

- **T-test (Welsh's t-test)**
- **U-test (Mann–Whitney test)**
- **Bootstrap test**
- **Kolmogorov-Smirnov test**

Reports

- **Time series plots (ensemble mean and spread for both ensembles)**
- **Spatial differences plots**
- **Matrix plot (see next slide)**

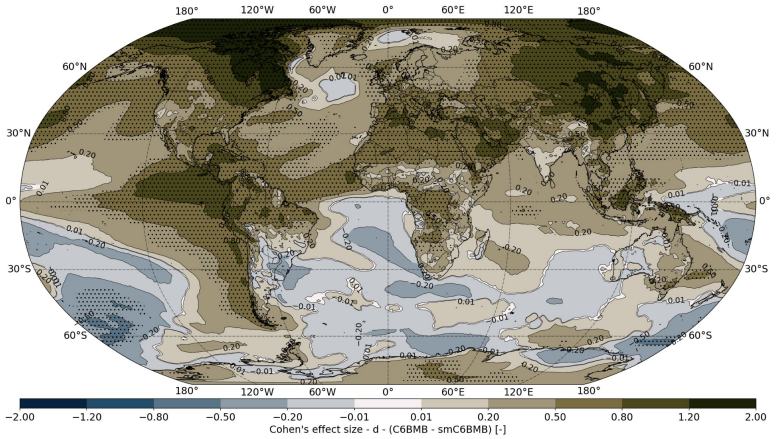

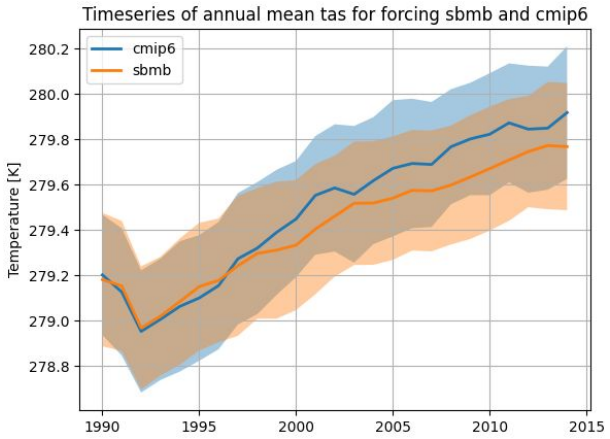

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Climate Model Replicability: Matrix Plot

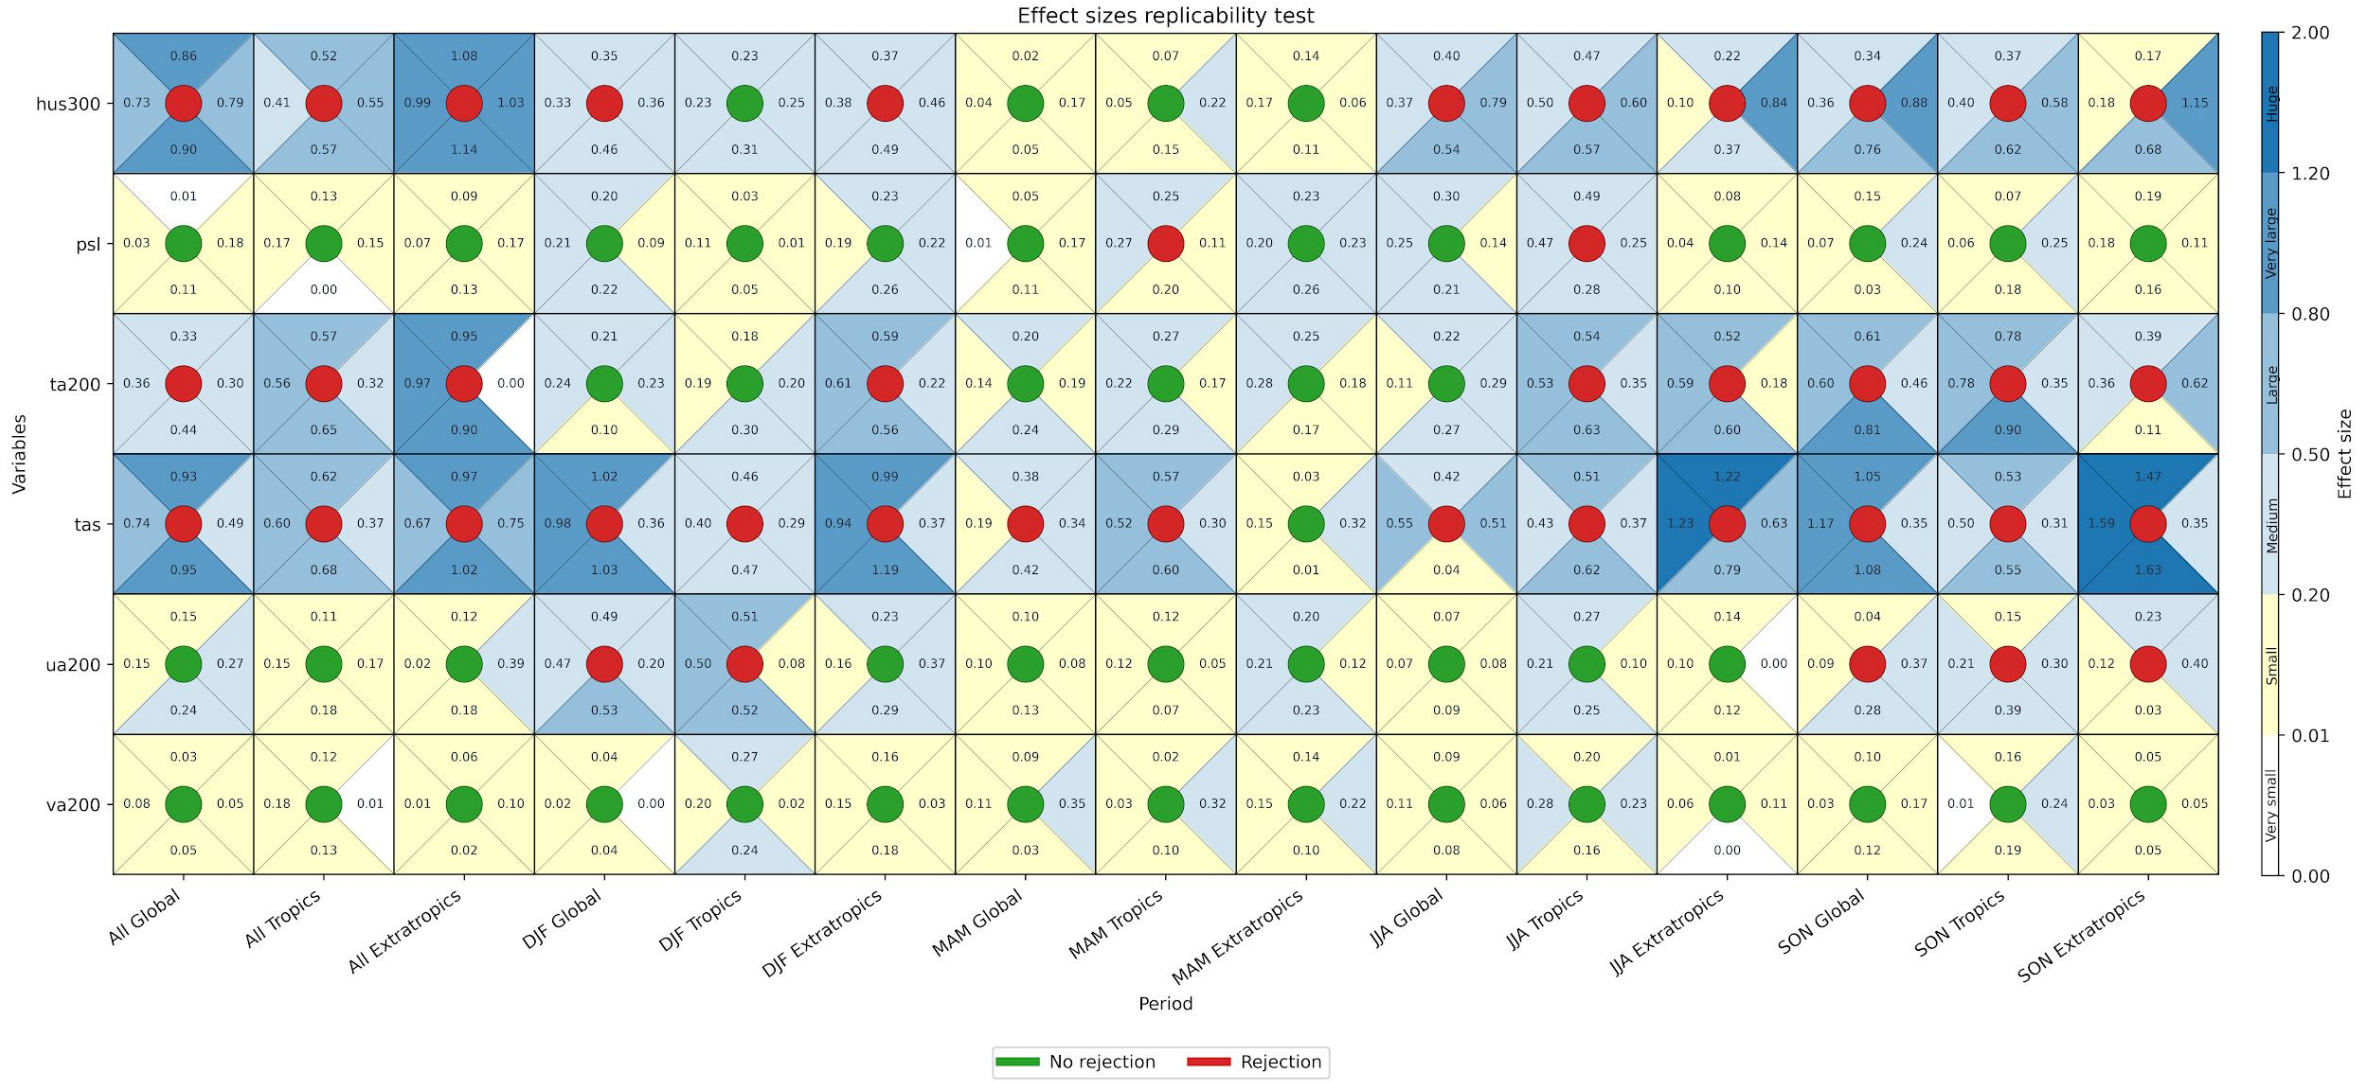

Projects

Work and proposals for the EC-Earth community around different projects

- ESiWACE Center of Excellence
  - Developing new profiling analysis methodology that can be applied to EC-Earth looking for new optimizations
  - HPC services using tools to facilitate M7 reduced precision
  - CPMIP improvement for easier integration and collection
- Horizon 2025-D1-01 and CMIP7
  - Proposed a WP to collect CPMIP: EC-Earth will be one key model that we want to support
  - Re-do our reproducibility exercise for CMIP7, proving the replicability and stability for our configurations
  - Other tools or activities as profiling could be useful to be applied in this context, please ask us for collaboration!

Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Thank you

Generalitat de Catalunya
Departament de Recerca i Universitats

Amb el suport del Departament de Recerca i Universitats de la Generalitat de Catalunya.
El grup "Computational Earth Science (CES)", Grup de Recerca de la Generalitat de Catalunya amb número d'expedient 2021 SGR 00785

xavier.yepes@bsc.es