esiwace

CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

ESiWACE3 - Excellence in Simulation of Weather and Climate in Europe, Phase 3

A Workflow for HPCW - The High Performance Climate & Weather Benchmark

Inpex, 15-17 April 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

A Workflow for HPCW - The High Performance Climate & Weather Benchmark

Mario Acosta & Miguel Castrillo
(mario.acosta@bsc.es, miguel.castrillo@bsc.es)

Barcelona Supercomputing Center (BSC)

Inpex, 15-17 April 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

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

Destination Earth

Destination Earth | EUROPEAN CENTRE FOR MEDIUM-RANGE WEATHER FORECASTS

HANAMI

HANAMI

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**

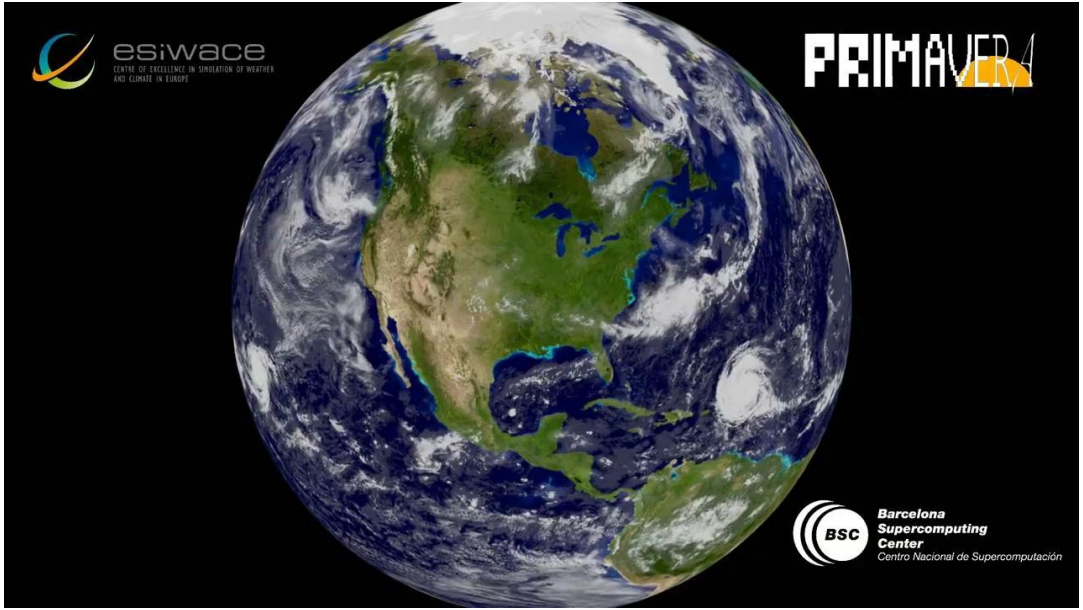

Ocean at different resolutions using the EC-Earth model performed by Oriol Tintó (BSC)

Coordinated BSC Barcelona Supercomputing Center Centro Nacional de Supercomputación

Consortium of 12 partners from 8 different countries

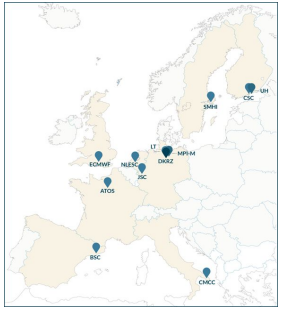

Start: 1 January 2023
End: 31 December 2026

ESiWACE3 ecosystem

Models/Tools

FVM

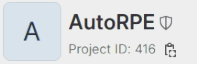

Field Compression Library

FDO4Climate

Auto Profile

Kernel Tuner Ecosystem

Field Compression Benchmark Suite

HPC Community

HPC3

ESM Community

HANAMI
HPC consortium EU - Japan

GANANA

HPCW - The High Performance Climate & Weather Benchmark

Why a Climate and Weather Benchmark?

HPCW - The High Performance Climate & Weather Benchmark

Why a new Benchmark?

High Performance Linpack

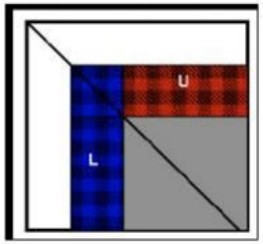

Dense linear equations
Still used for Top500 ranking (and Green500)

1979

STREAM

STREAM
High Performance Computing

Pure memory benchmark

1996

High Performance Conjugate Gradients

HPCG
High Performance Conjugate Gradients

Balanced memory/compute benchmark

2014

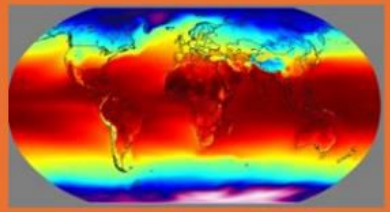

Benchmark adapted for weather & climate considerations ?

Slide provided by Erwan Raffin, Atos

HPCW - The High Performance Climate & Weather Benchmark

Why a new Benchmark?

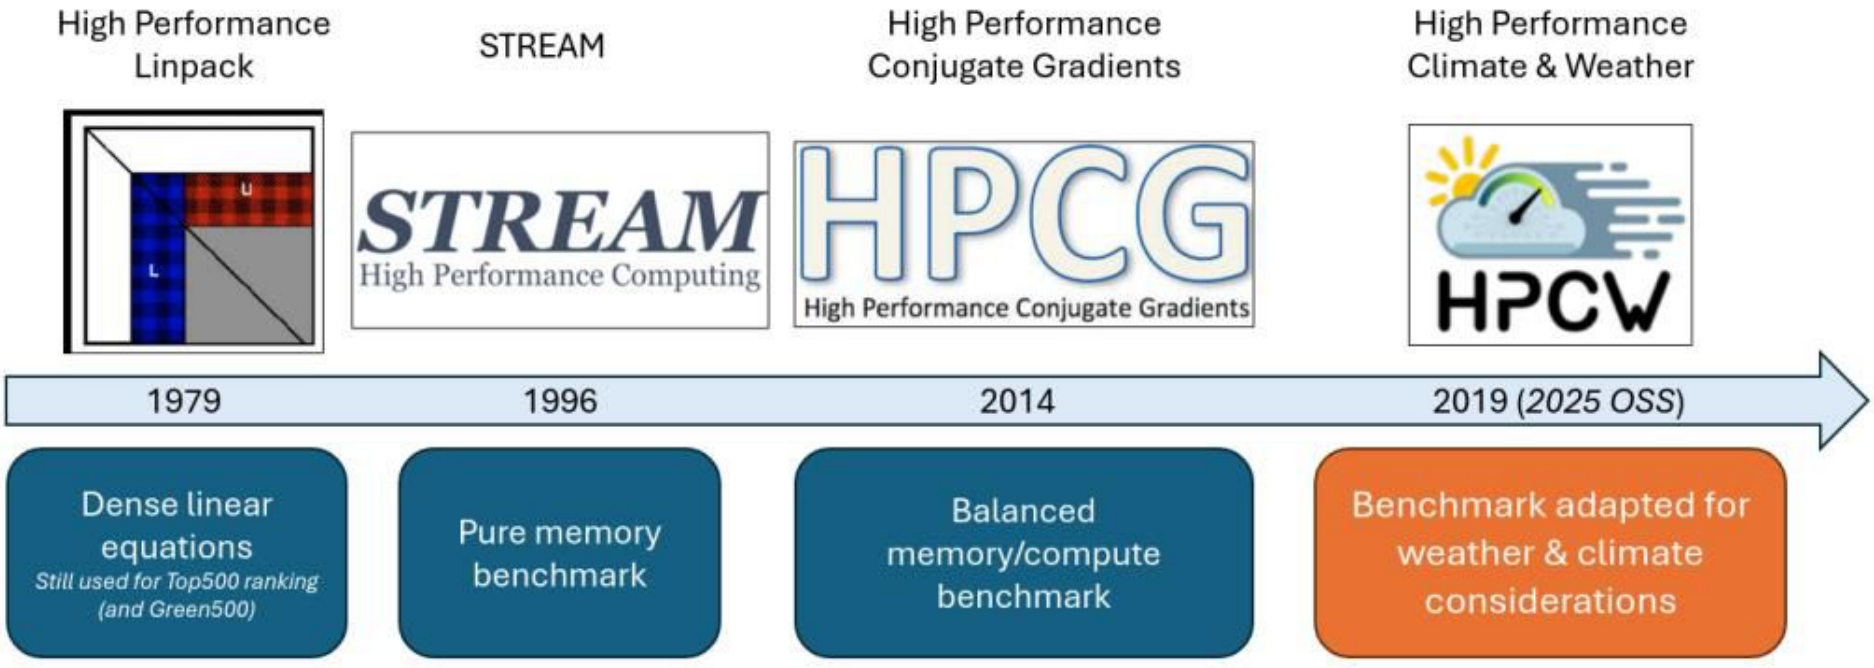

Slide provided by Erwan Raffin, Atos

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.)

Slide provided by Erwan Raffin, Atos

How do HPCW's technical choices promote reproducibility?

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 (in files described later on)
  - stored in the Git repository as well.

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 (in files described later on)
  - 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

Slide provided by Erwan Raffin, Atos

HPCW Roadmap In Europe and beyond

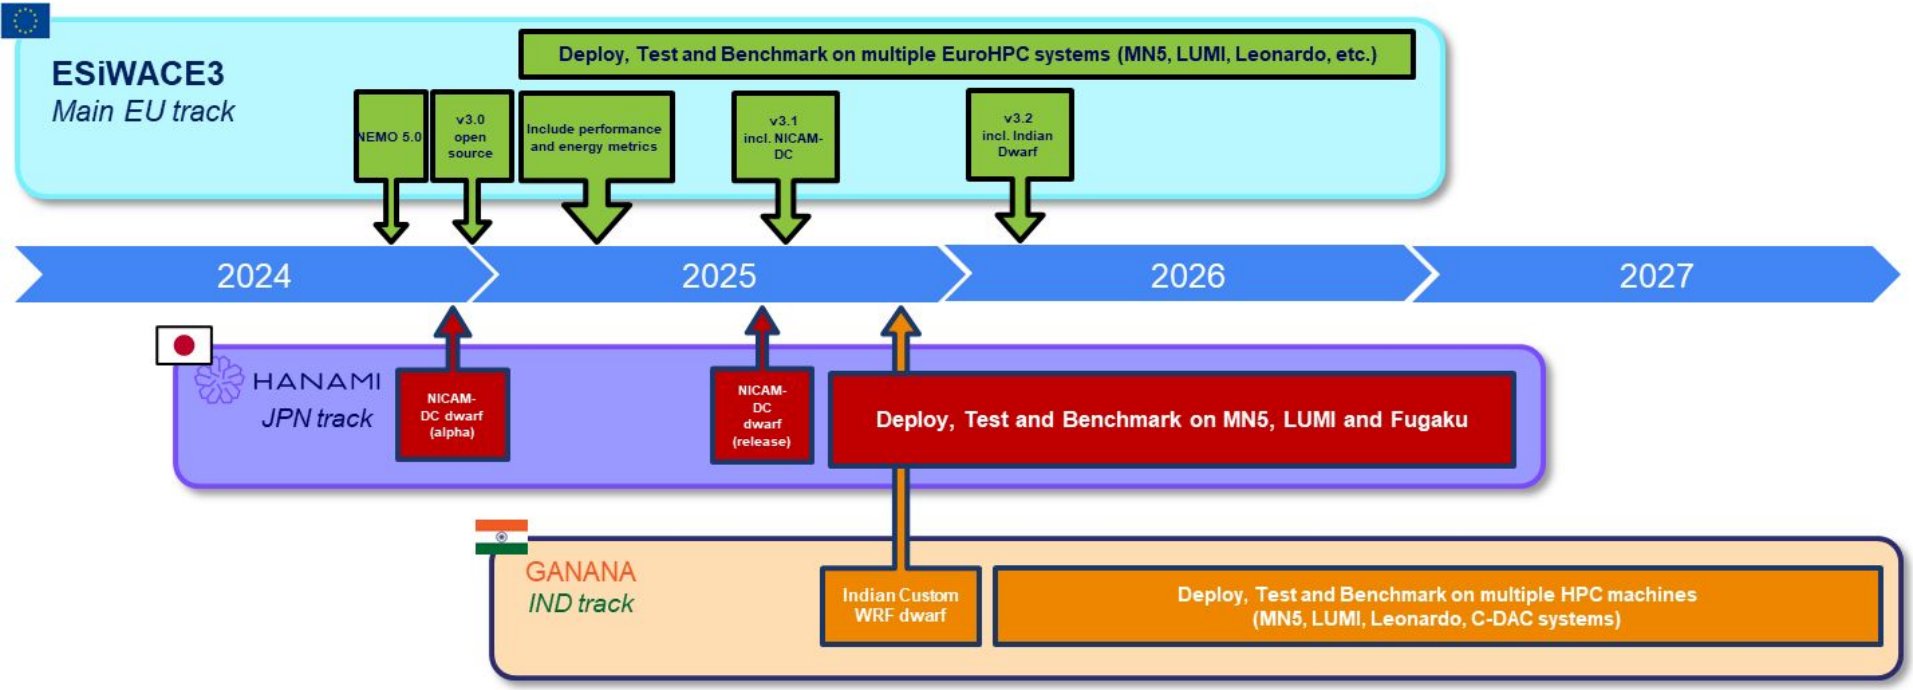

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

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

AS API

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 GUI

BSC 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

Integrating communities, platforms and standards

Destination Earth

European Digital Twin Ocean
EDITO Model Lab

Digital Twins

Workflow Run RO-Crate

METACLIP

Provenance

ecFlow

AiiDA

LEXIS 2 Platform

Workflow managers

jupyter

OpenID

GitHub Oauth Apps

Interfaces

HPCs: BSC/MareNostrum, CSC/LUMI, CINECA/Leonardo, RIKEN/Fugaku, ECMWF/HPC2020, LuxProvide/Meluxina

AUTOSUBMIT

Challenges and roadblockers

1. **Access to HPC resources** (as a community) [HEP, SKA, **ESiWACE**]
2. **Interface to federation of resources** (e.g. EuroHPC) [HEP, **ESiWACE**]
3. *Co-design for the post-exascale systems* [HEP]
4. Portability of the benchmarks and codes [ESiWACE, SKA]
5. Deployement issues of the workflows and applications [ESiWACE, SKA]
6. Metrics for different deployement scenarios [ESiWACE]
7. Semantics and quality of data [Urgent computing]
8. Resources provisionning [Urgent computing]
9. *End-to-end workflow control* [Urgent computing]
10. Multi-tenancies [Urgent computing]
11. Data logistics [SKA]

Specific challenges

- Climate experiments for IPCC reports (such as CMIP) require community efforts with common (huge) experiments → HPCW need access and resources for different platforms and a lot of computational hours
- Federated access require preparing the workflow manager to take care of different user configurations for the particularities of each platform (job system, storage management, Computational heterogeneity, etc).
    - A multi-platform workflow manager such as Autosubmit will allows resources to be federated under a single workflow configuration running the same software on different types of platforms.
- Automatization of the deployment and execution process for huge experiments and different hardware

Challenges and roadblockers

1. Access to HPC resources (as a community) [HEP, SKA, ESiWACE]
2. Interface to federation of resources (e.g. EuroHPC) [HEP, ESiWACE]
3. ***Co-design for the post-exascale systems*** [HEP, **ESiWACE**]
4. Portability of the benchmarks and codes [ESiWACE, SKA]
5. Deployement issues of the workflows and applications [ESiWACE, SKA]
6. Metrics for different deployement scenarios [ESiWACE]
7. Semantics and quality of data [Urgent computing]
8. Resources provisionning [Urgent computing]
9. *End-to-end workflow control* [Urgent computing]
10. Multi-tenancies [Urgent computing]
11. Data logistics [SKA]

Specific challenges

- Use in an efficient way HPCW to improve the Fortran compiler of emergent technologies
- Facilitate the co-design to improve the compiler optimizations for new accelerators in the post-exascale context
- Use HPCW for co-desing and improvement of the hardware for the better adaptation of Earth System Models for more efficient climate change adaptation

Challenges and roadblockers

1. Access to HPC resources (as a community) [HEP, SKA, ESiWACE]
2. Interface to federation of resources (e.g. EuroHPC) [HEP, ESiWACE]
3. *Co-design for the post-exascale systems* [HEP, ESiWACE]
4. **Portability of the benchmarks and codes** [ESiWACE, SKA]
5. **Deployement issues of the workflows and applications** [ESiWACE, SKA]
6. **Metrics for different deployement scenarios** [ESiWACE]
7. Semantics and quality of data [Urgent computing]
8. Resources provisionning [Urgent computing]
9. *End-to-end workflow control* [Urgent computing]
10. Multi-tenancies [Urgent computing]
11. Data logistics [SKA]

Specific challenges

- Deploy and port HPCW seamlessly in different platforms
- Automate the deployment and operation of HPCW using an automated workflow
- Integrate the workflow with a workflow manager (Autosubmit) that can act on-demand
- Ensure that the workflow report and the workflow manager collects the relevant performance metrics
    - Design and implement new performance metrics
- Test the HPCW-based workflow on different platforms and compare the results
    - Validate the portability and on-demand capabilities (Digital Twins oriented)

Challenges and roadblockers

1. Access to HPC resources (as a community) [HEP, SKA, ESiWACE]
2. Interface to federation of resources (e.g. EuroHPC) [HEP, ESiWACE]
3. *Co-design for the post-exascale systems* [HEP, ESiWACE]
4. Portability of the benchmarks and codes [ESiWACE, SKA]
5. Deployement issues of the workflows and applications [ESiWACE, SKA]
6. Metrics for different deployement scenarios [ESiWACE]
7. Semantics and quality of data [Urgent computing]
8. Resources provisionning [Urgent computing]
9. ***End-to-end workflow control*** [Urgent computing, **ESiWACE**]
10. Multi-tenancies [Urgent computing]
11. Data logistics [SKA]

Web

Cloud

HPC

Specific roadblocks

- Portability and interoperability
- Validation: end-to-end performance metrics
- Representativeness of the benchmark
- Communication with vendors and hardware developers
- Integrate our benchmark in CI/CD frameworks
- Opportunities to participate in the federation of the HPCW workflow
- Opportunities to provide our benchmark for testing novel hardware for post-exascale supercomputing

THANK YOU!

hanami@saxifrage.saclay.cea.fr

FOLLOW US

Funded by the European Union

EuroHPC Joint Undertaking

This project received funding from the European High Performance Computing Joint Undertaking (EuroHPC JU) under the European Union's Horizon Europe framework program for research and innovation and Grant Agreement No. 101136269. Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or EuroHPC Joint Undertaking. Neither the European Union nor the granting authority can be held responsible for them.

esiwace

CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

Interested in getting in touch?

Website: www.esiwace.eu

YouTube: https://www.youtube.com/@esiwace880

X: https://x.com/esiwace

Bluesky: https://bsky.app/profile/esiwace-community.bsky.social

LinkedIn: https://www.linkedin.com/company/esiwace3

zenodo

ESiWACE is on Zenodo, the Open Access repository for scientific results
https://zenodo.org/communities/esiwace

OPEN ACCESS

mario.acosta@bsc.es

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

Breaking barriers for decision making

Some limitations for the use of environmental data in different socio-economic sectors.

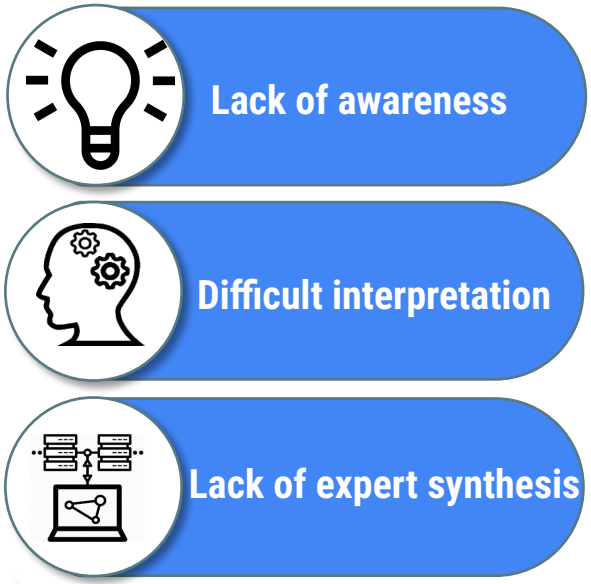

Possible solution: to distil the information from existing sources to be integrated in decision-making.

Environmental services

- **Goal:** the development and incorporation of environmental data for planning, policy-making and practice at the global, regional and national scale.
- **Implementation method:** co-production and co-design.

DESTINATION EARTH INITIATIVE

- **Context**: European Commission's programme, part of **Green Deal & Digital Strategy**.
- **Objective**: To develop **digital twins** (DTs) of the **Earth** to support decision-making.
- **Users**: Policy makers and environmentally sensitive sectors.
- **First DTs**: Climate Change Adaptation Digital Twin and Weather Extremes Digital Twin.
- **Computing resources**: Two EuroHPC pre-exascale supercomputers: LUMI (currently in operation in Finland) and MareNostrum 5 (Spain) provided by EuroHPC Joint Undertaking.

CSC
LUMI

BSC Barcelona Supercomputing Center
Centro Nacional de Supercomputación

MareNostrum 5

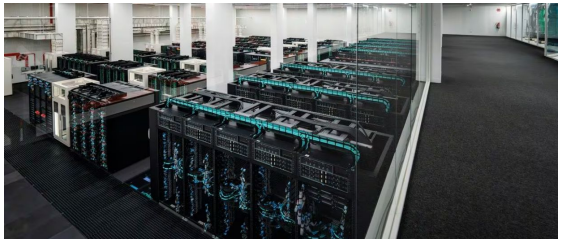

Climate Digital Twin

- Global climate simulations at **unprecedented horizontal resolution.**
- **User-driven approach** focused on user interactivity.
- Novel approach with **streaming of climate model output** to impact models.
- Designed to **support decision-making** on the impacts of climate change and different adaptation strategies.

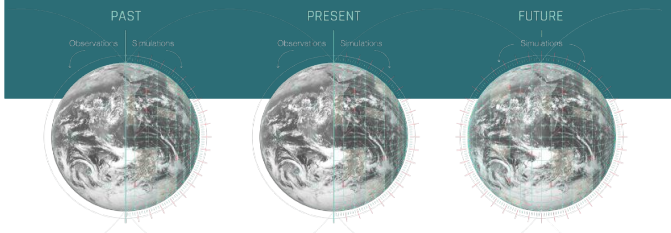

Extremes Digital Twin

- Short-range predictions with and **increased resolution** to sub-km scales.
- Address the **impacts of extreme events**
- Forecast of extreme **hydrological** events.
- Simulate hazardous **air quality** situations.
- Anticipate extreme events impacting the **energy** production.

Climate Dig … Digital Twin

- Global clima… predictions with and **unprecedent…** **…esolution** to sub-km scales.
- **User-driven …** **…impacts of extreme events** interactivity.
- Novel a… …recast of extreme **hydrological** events. climate…
- Designe… …ulate hazardous **air quality** the imp… …ations. differen…
- …cipate extreme events impacting the …**rgy** production.

Earth Sciences Department
BSC Barcelona Supercomputing Center
Centro Nacional de Supercomputación

Our contribution is fundamental for the Climate Digital Twin

We lead important aspects such as:

- Portability, scalability, development and validation of the models
- Workflow deployment
- Data flow
- Use cases