

European objectives for Models, Tools and HPC in the future research infrastructure

The mission and objectives of the European Network for Earth System Modelling (ENES) for Models, Tools and HPC services.

Authors: Bryan Lawrence (NCAS), Fanny Adloff (DKRZ), Jean-Claude André (Cerfacs), Joachim Biercamp (DKRZ), Uwe Fladrich (SMHI), Adrian Hines (STFC), Sylvie Joussaume (IPSL), Martin Jukes (STFC), Kim Serradell (BSC), Sophie Valcke (Cerfacs)

Introduction

The European Network for Earth System Modelling (ENES) has operated since 2001 to support advances in climate modelling and understanding¹. ENES has been supported by a succession of infrastructure projects (IS-ENES, IS-ENES2, and IS-ENES3; 2009-2022) as well as by an HPC Centre of Excellence (ESiWACE, ESiWACE2; 2015-2022). These activities, carried across multiple institutions exploiting mutual support, have been delivering on a community infrastructure strategy first published in 2012 and updated in 2017², and jointly deliver the ENES Research Infrastructure (ENES-RI).

The “Models, Tools and HPC” part of the ENES-RI addresses the simulation part of the Earth system modelling workflow by gathering and addressing community requirements, supporting and accelerating common software development, facilitating the efficient use of resources, preparing for next generation computing, and upskilling the community with networking and training opportunities. Together the services required for these activities target community codes, software optimisation, optimal hardware utilisation, and community expertise.

This document outlines the overall objectives of the “Models, Tools and HPC” activities, with some details in terms of current activities, and near-term aspirations. It will be updated as the community infrastructure strategy evolves.

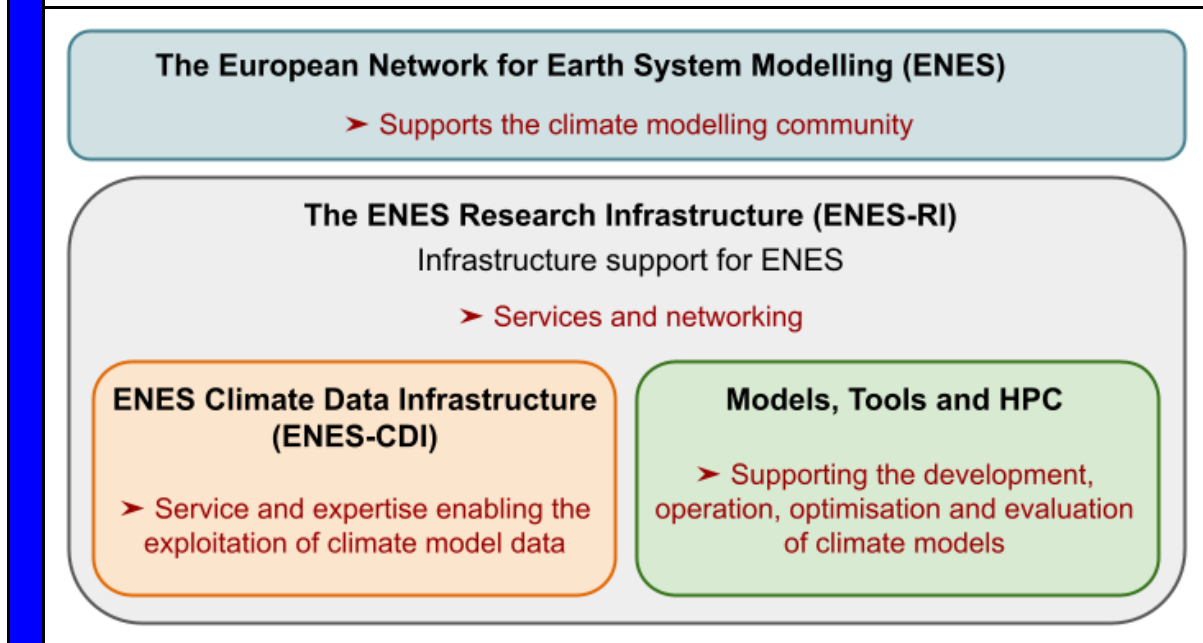


IS-ENES3 & ESiWACE2 projects have received funding from the European Union's Horizon 2020 research and innovation programme under the grant agreements No 824084 & 823988

¹ <https://portal.enes.org/about-enes>

² <https://portal.enes.org/strategy/>

Figure 1: Placing “Models, Tools and HPC” in the ENES ecosystem; one component of the research infrastructure created by the ENES community.



General Objectives

Supporting and Accelerating of Model and Tool Development

Earth system models include a range of components and tools, many of which are shared across modelling systems. The ENES-RI provides services to coordinate the necessary shared development, promote uptake and common tooling, and where possible invests in accelerating the development of both model code and analysis tools.

Facilitating the Efficient Use of Resources

Earth system models use large amounts of computing and storage resources. The ENES-RI assists the community in the efficient use of these platforms - and the necessary networks and data systems - by coordinating requirements, sharing best practice, providing a single point of contact for organising shared access, and facilitating code optimisation.

Preparing for Next Generation Computing

Computer systems are changing rapidly, with new chip technologies and new system architectures leading to considerable hardware heterogeneity, but most existing software cannot easily exploit those new capabilities. The ENES-RI facilitates the exploitation of new computer platforms by fostering the development of new science code and middleware to make the best use of new computing and storage hardware.

Nurturing the Community

Model development requires scientists and engineers with specialised skills which are not easy to gain in small groups. The ENES-RI facilitates networking between individuals and groups, increasing the opportunities for training at all levels, with the aim of increasing the number of such people in Europe.

Specific Activities

A number of specific activities are required to deliver the general objectives and sustain the cooperation necessary to develop future model and data technology.

Accelerating Progress by fostering collaboration

Modelling is used to improve understanding of the climate system and to predict and project the future. Advancing such modelling depends on improving the accuracy and representativeness of models, on understanding variability, and on ensuring models can simulate the past and important aspects of future scenarios. Improving accuracy and representativeness is context dependent, requiring appropriate combinations of resolution and complexity and appropriate comparison tools and data. Many simulations may be necessary to understand variability and examine future scenarios. At the frontiers of current science, most of these activities cannot be carried out by lone individuals or even single institutions. ENES exists to foster the necessary community collaboration on model development, to enable the efficient use of community resources for using these models to carry out large-scale simulations, and to support the analysis and use of the data produced.

Representing the modelling community

Community information sharing, co-development of software and effective use of large national and pan-European computing platforms, all require a blend of communication and representation. ENES provides venues for communication and delivers a representative task force of community experts who can reach out across borders and between institutions to foster the efficient use of available resources; to represent the weather and climate community, and strengthen collaboration with industry, EuroHPC, other HPC stakeholders and other European actors providing services to, or using services from the climate community.

Advancing High Resolution Modelling

An important aspect of model fidelity is resolution. The ENES-RI supports the European community in the development and execution of the highest possible resolution of global weather and climate models. Activities include coordinated campaigns to push resolution and scalability as well as the development and maintenance of tools to mitigate, manage, and exploit the consequential data deluge.

Shared Model and Tool Components

ENES encourages the sharing of tools and model components within the community to address both knowledge exchange and the avoidance of duplicated effort. Where it is necessary to improve sustainability and/or add critical mass ENES facilitates support for community projects such as couplers, workflow engines and schedulers, IO servers and shared science codes such as (currently) the European sea ice and ocean models. One example of shared activity is the Earth System Model Evaluation Tool (ESMValTool), a community tool for capturing the diagnostic and performance metrics needed for the evaluation of earth system models. As with the ocean and sea ice communities, ESMValTool has its own development plan, the ENES role is to foster additional support and where necessary facilitate governance.

Fostering Flexible Software

The computing ecosystem is changing rapidly with significant heterogeneity after a long period of homogeneous systems. With hardware life cycles typically measured in years and software life cycles measured in decades maintaining efficient codes is becoming problematic. ENES supports the community in exploring the development of middleware to abstract the scientific parts of the code from the underlying hardware in order to facilitate model codes which can remain performant across rapidly changing hardware solutions.

Services to support policy relevant science

The delivery of operational climate services and direct science advice will be done by others often using data from coordinated experiments. ENES contributes to the coordination of international programmes, including those of the World Climate Research Programme, which involve reference experiments and the production of international standards for modelling and model-to-data-to-information workflows.

Supporting Access to Model Codes

One of the central tenets of climate science is the need for model diversity, and for models to be used in the community as well as in large modelling centres, however ESM systems are technically difficult to deploy and use. ENES promulgates common services for access to model codes and documentation, encourages permissive licences, and provides a directory of accessible model codes along with information as to how they may be accessed and used.

Training Opportunities

The weather and climate community recognise that productive research builds on well engineered codes which are easy to understand, extend, and run. Model developers are often software engineers without much modelling experience, or scientists without much formal software development training. The ENES-RI instigates, advertises, and coordinates targeted training opportunities wherever possible. Where funding allows, ENES actively delivers events such as summer schools, workshops, coding sprints and other mechanisms to upskill the community.

Porting and Optimising Codes

Time-to-solution and efficient use of HPC require code optimised for a given platform, and with the need to use multiple machines, also means a requirement for performance portability. Coding for performance, and porting the resulting codes requires advanced technical skills. Within the wider training objectives, ENES facilitates the provision of support and training for code profiling and optimisation, promotes coding standards, and fosters the use of open and common performance tools.

Standardised Benchmarking

The procurement of new computing systems which are fit for purpose is a heavy burden on the community, as is ensuring that next generation software will work on next generation hardware. Benchmarking is a key tool in both developing software and procuring systems, and the use of such benchmarks is a key tool in the co-design of software and hardware environments. ENES encourages the development of community benchmarks, standardised datasets (e.g for machine learning), and all the necessary documentation; and advertises their availability.

Technology Tracking

Technology is changing rapidly. Along with a massive proliferation of new hardware for computing and storing data, new mathematical and numerical techniques are being designed to exploit massively parallel accelerated machines. Machine learning and Artificial Intelligence techniques are being introduced into both simulation codes and analysis workflows. ENES organises and, where possible delivers, community networking activities to exploit information from other disciplines and share within the earth system simulation community.