




Call for Proposals

Collaborations to prepare Europe's Weather and Climate Models for pre-exascale systems (ESiWACE-S1)

2019-2020

A joint call for proposals by the Netherlands eScience Center and Atos-Bull



Part of the Services provided by the ESiWACE2 Center of Excellence
in Simulation of Weather and Climate in Europe



The projects ESiWACE and ESiWACE2 have received funding from the European Union's Horizon 2020 Research and Innovation Programme under Grant Agreements No 675191 and 823988.



Contents

Introduction	3
Background	3
Purpose of this call	4
Available budget	5
Validity of this call	5
Guidelines for applicants	6
Who can apply?	6
What can be applied for?	6
When can applications be submitted?	6
Preparing an application	7
Specific conditions	7
Submitting an application	7
Assessment procedure	8
Procedure	8
Admissibility and assessment criteria	8
Contact details	10
Contact	10
Other information	10



1 Introduction

1.1 Background

Weather and climate models are large and complex applications that experience a tension field between investments to enhance certain features, increasing fidelity, spatio-temporal resolutions, or resolving more physical processes, and investments to adapt the software to the latest hardware architectures. Generally, these improvements go hand in hand, as higher accuracy often only can be achieved through increased resolution or extending the ensemble size, and thus increasing the required computational resources. Where the modeling developments are driven by scientific expertise, the performance optimization is driven by the desire to use the compute resources efficiently and at sufficient scale.

However, after a long period of relative stability, the computing infrastructure is rapidly evolving and diversifying, with more change on the horizon. Adapting existing software to the latest computing platforms and exploiting the full capability of the architecture requires specific expertise and extensive changes at the source code level. As such, continued development of weather and climate models requires new collaborations between experts from different fields.

This call for proposals aims to support the exascale preparations for the weather and climate modelling community in Europe. We create short collaboration projects that provide guidance, engineering, and advice to improve model efficiency and port models to existing and upcoming computing infrastructure. All groups developing and maintaining weather and climate codes - not only the ESIWACE2 partners - can apply. Proposals will be subject to a review regarding technical feasibility by the involved partners, being the Netherlands eScience Center and Atos-Bull, and undergo a review more focussed on scientific impact by an external committee. When found eligible, the project will be granted in-kind contribution by one of these partners.

The Netherlands eScience Center is the Dutch national center of excellence for the development and application of research software to advance academic research. The eScience center employs a large team of eScience Research Engineers, who are experts in state-of-the-art computational methods and technologies and have a keen interest and experience in developing research software. The eScience center is involved in more than 90 collaborative projects, spanning many different research disciplines and application domains, of which 11 currently projects are within the domain of weather and climate. The center harbours unique expertise on programming for heterogeneous hardware architectures, optimizing software for supercomputing platforms and implementing new features and technologies in the often extensive fortran codes that govern the domain of atmospheric physics and oceanography.

The Atos-Bull team in charge of the ESIWACE services, part of Applications and Performance team, is composed of experienced HPC experts involved in collaborative European projects, such as ESCAPE and ESIWACE in the weather and climate domain, but also in other domains as VECMA (Verification Validation and Uncertainty Quantification), ASPIDE (Exascale programming models for extreme data processing), ChESEE (Center of Excellence in Solid Earth), CompBioMed (Center of Excellence in Computational Biomedicine). The core expertise of the team is HPC applications profiling and tuning and a variety of high-end HPC systems including heterogeneous systems, memory accelerator, etc. Profiling



and optimization can be done at different levels, from processor microarchitecture to many-node systems, to take advantage of its characteristics and thus reaching the best performance and scalability. Thanks to Atos-Bull's position as European leader in HPC, its Research and Development department and its strong relationship with technology providers as well as with the weather and climate community, the team has access to a wide set of state-of-the-art HPC hardware, software, tools and knowledge that allow it to help at preparing weather and climate applications for exascale HPC systems.

1.2 Purpose of this call

This call for proposals aims to support the weather and climate model developing communities in Europe with preparing their simulation codes for (pre-)exascale computing platforms. The goal is to improve model efficiency and readying the software to enable model execution on existing and near-future hardware architectures, and as such, enable simulation experiments at unprecedented grid resolutions or ensemble sizes, or include computationally expensive physical processes that were previously unfeasible.

The goal of this call is to create collaborations that provide consultancy and engineering efforts to support the exascale preparations of weather and climate model developing communities. These collaborative projects allow experts in HPC and accelerated computing to work together with the model developers to advance the codebase such that (parts of) the model can be executed efficiently on modern CPU processor or modern computing accelerators, such as Graphics Processing Units (GPUs). Achieving such an ambitious goal in a short-lived project requires very active involvement from both sides. We expect the development team to provide a model configuration that is widely adopted in the model user community, presenting a code execution path that deserves being accelerated. In the same spirit, we expect guidance on the typical use cases, testing procedures, coding style standards, library dependencies and version control system to smoothen the collaborative development process and ensure correctness of our results.

A competitive proposal aims to have long-lasting impact on (the main branch of) the source code of an established atmospheric or oceanographic model (or any other climate model component) to progress its capability of efficiently and effectively using modern computing hardware. To ensure such impact, there is a strong preference for models that are developed under sustainable and re-usable software engineering practices and are developed as open source projects with a permissive open source license.

We encourage proposals to target accelerator hardware such as GPUs, which will most likely constitute the bulk of the capability of future exascale systems. In such cases we will aim at porting the critical path of the code to the accelerator and tuning the resulting code for maximal throughput on the device. To achieve this in a sustainable fashion, parts of the code may be rewritten using domain-specific languages (DSLs) or stencil libraries with CUDA- or OpenCL-backends.

Proposals for this call may also target traditional CPU architectures or many-core variants. In such cases, the effort may be focussed on e.g. improving cache efficiency, stimulating the use of SIMD instructions, increasing floating-point instruction rate by moving to lower precision numerical representation,



implementing fine-grained or task-level parallelism, detection and elimination of MPI bottlenecks or reducing the pressure on the storage system.

It is encouraged, but not a necessary condition, that the applicant gives the Research Software Engineers (RSEs) access to the target platform for development and testing. If this is impossible, it should be discussed how to gain access to the target hardware through the ESIWACE2 PRACE allocation or one of Atos-Bull's facilities. Finally we note that the specifics of future exascale systems remain unknown, and that preparing models for future exascale platforms is currently a moving target.

1.3 Available budget

In this call, each grant consists of in-kind support in the form of one or two skilled Research Software Engineers (RSEs) employed by either the Netherlands eScience Center or Atos-Bull. The RSEs will work remotely on the project for 6 person months (PMs) during the year 2020. The hours will be planned flexibly and in mutual agreement with the applicant. Included in this time is optionally a visit to the modeling group developing the source code and the dissemination of the work, both internally in e.g. ESIWACE project meetings and externally through conference and workshop presentations. In the years 2020 and 2021, new calls for proposals will be issued targeting projects in respectively 2021 and 2022. Applicants that have applied in earlier calls can apply again in later calls.

Given the total budget, four proposals can be granted as part of this call. After submission a technical review is carried out focussing technical feasibility and state-of-the art technology. The proposals that have passed the technical review are sent to the external scientific review committee for review, where the final selection is made based on the scientific merit and impact. Finally, in December 2019 the grant decisions will be announced and the applicants of awarded proposals will be contacted.

1.4 Validity of this call

This call is valid for proposals submitted before **November 1st 2019, 14:00 CET**. The call is open from **September 6th 2019** until **November 1st 2019, 14:00 CET**: proposals sent after this time will no longer be considered for grants in 2020.



2 Guidelines for applicants

2.1 Who can apply?

This call addresses all European model developing groups in Weather and Climate, including atmospheric sciences, oceanography, and climate-related domains such as land or sea ice modeling, land-surface modeling, atmosphere or ocean (bio-)chemistry, etc. The call includes, but is not restricted to global models, and therefore regional models or small-scale very high-resolution applications such as atmospheric large-eddy simulation codes or morphological coastal models are targeted too. Members of the ESIWACE consortium are allowed to apply for this call, but groups outside of ESIWACE have a strong preference.

We aim to have a permanent impact on the source code of the application, and this is most likely achieved if the applicant is (a member of) the core development team of the model code; derived 'downstream' versions or user-specific branches will have a smaller chance of being granted, as we wish not to contribute to the fragmentation of the European earth system modeling landscape.

2.2 What can be applied for?

The grant consists of 6 PMs of in kind consultancy and guidance by one of the expert RSEs on high-performance and heterogeneous computing by either NLeSC or Atos-Bull. These engineers will work alone or in pairs on the project and fulfill the hours throughout 2020. Although the engineers will work predominantly remotely, we expect a matching effort from the granted applicant's research group to establish a smooth collaboration with sufficient support for the RSE. During the project there will be at least one visit of the RSEs to the applicant's research facility.

Because the work happens on a project basis, there is no available budget for long-term maintenance of the contributed code. The long-term maintenance will be on the applicant's modeling group. To assist with this, a technical report will be written by the RSE and disseminated to the model development group. Furthermore, the contributed code will be documented according to standard good practice, or the modeling group internal coding standards.

Finally, the work will be disseminated in conferences and workshops by the RSEs, offering exposure for the applicant and promotion of the software. It is encouraged that the applicant and the RSEs together promote and publish about the work carried out in the project. These efforts are considered part of the collaboration, and the time for RSEs to spend on these activities is part of the in-kind budget of the grant.

2.3 When can applications be submitted?

Proposals can be submitted in the period **September 6th 2019** to **November 1st 2019, 14:00 CET**.

Further information about the procedure is found in Section 3.1.



2.4 Preparing an application

A proposal can be submitted by completing a google form within the above time frame. The form consists of three sections addressing the following requested information:

1. A brief overview of the model, its scientific impact and societal relevance. An explanation and underlying motivation of the deployment of the model at the exascale.
2. A brief description of the targeted platform or hardware architecture. The limiting factors to the large-scale deployment of the application onto the target platform and some ideas about what performance gains can be expected (the latter may be somewhat speculative). A list of components or routines that need to be ported or optimized to achieve this speedup.
3. Plans for the sustainability, dissemination and maintenance of the developed software.

Applications must be completed in English.

2.5 Specific conditions

We expect the applicant to either provide or adopt basic infrastructure for (remote) software development such as a version control system, a platform to keep track of issues, progress and discussions (e.g. jira, github, ...) and to provide a representative test case for monitoring numerical results and performance of the program. In general, the applicant should display an open, collaborative mindset and is expected to give feedback to the RSEs during the development cycle. This means that the applicant is responsive to emails and announces periods of absence ahead of time.

The target hardware platform, together with all necessary compilers, firmware and performance measurement tools, has to be available to the engineer during the development process. Access to the target platform can be established either through (i) a joint external (pilot) project call to a European facility, (ii) the internal PRACE allocation of the ESIWACE2 project, or (iii) permission for some internal facility given by the applicant to the RSEs. In the latter case, it is important that the allocation of the resources (nr. of nodes, disk space) and permissions to install the required tools is sufficient for running full-scale tests and we expect a good collaboration with the local system administrators.

During any visits of the RSEs to the applicant's institute, the applicant is responsible for providing the RSE with a suitable, safe, and healthy working environment. This includes providing an environment that is harassment free and free of unwanted behavior, for more information see [code of conduct](#).

When the above conditions are systematically violated, the ESIWACE2 consortium may choose at any moment to stop the project funding. In particular, the first two months of projects are considered to be a trial period, in which the effectiveness of the collaboration is evaluated and if deemed not effective the project can be terminated.

2.6 Submitting an application

The online application form will be made available on <https://www.esiwace.eu/services/hpc-userservices>.



3 Assessment procedure

3.1 Procedure

First there will be a technical review of the proposal by either the engineering team at NLeSC or Atos-Bull. This review will not assess the scientific impact, but rather focus on the technical feasibility of the project: is the targeted performance improvement realistic, is the code practically portable to a new architecture within the timeline of the project, and can the underlying numerical algorithm be accelerated? We reserve the option to invite the applicant for a remote consultation meeting as part of the technical review process. The outcome of the technical review is binding, only a selection of applications will pass to the second round for scientific review.

The final decision on the grant will be made by an external scientific review committee of four renowned scientists in the fields of weather and climate, with affinity for HPC, two of which are from within the ESIWACE2 project and two are external reviewers.

In December 2019, the proposals will have been reviewed and the granted projects will be announced. After that, an RSE with matching expertise from either the Netherlands eScience Center or Atos-Bull will be appointed to the project and a project planning will be made in early 2020.

3.2 Admissibility and assessment criteria

A proposal is considered admissible if the applicant is affiliated with a European non-commercial research institution, fulfills the description of 2.1, and there is no conflict of interests between the proposed project and the service providers. Furthermore, the code needs to have a license and the first-round reviewers (and eventual RSEs) need to be able to access to the source code as part of the review process. Finally, the intended hardware platform has to exist within Europe, and obtaining access by the engineers to the hardware should be plausible.

The first round technical review will focus on the following criteria:

1. Technical feasibility: are the proposed code changes doable within the 6pm time frame? Is it plausible that the resulting code will run faster on the proposed target platform?
2. Technological state-of-the-art: are the applicants aware of the current state-of-the-art algorithms and techniques in their domain? Is the target hardware likely to be part of future European exascale machines and does the proposed work bring the application closer towards exascale deployment?
3. Software sustainability and lateral impact: does the proposal provide a credible pathway towards sustainable software? Will the proposed changes be merged to the main branch of the model, and is this code sufficiently maintained by a community?

The second, decisive review round focuses on the scientific impact of the project, such as

1. Scientific quality: does the proposed work enhance the quality of the model output, does the increased performance of the model allow its users to reach unprecedented accuracy or enable scientific breakthroughs?



2. Scientific novelty: does the proposed work allow the model community to explore new territory, use the software in other scientific applications?
3. Lateral impact: does the proposed work have scientific impact in other areas or communities? Are there possibilities for generalization of the contributed software?



4 Contact details

4.1 Contact

If you have specific questions about this call for proposals and the assessment procedure, please contact:

Open calls at eScience center:

Email: open-calls@esciencecenter.nl

For questions about the ESIWACE2 Centre of Excellence in Simulation of Weather and Climate in Europe, please contact:

ESiWACE2 Consortium

Email: esiwace@dkrz.de

For questions about the Netherlands eScience Center in the context of this call, please contact:

Dr. Ben van Werkhoven, eScience Research Coordinator

Tel.: +31 20 4604770

Email: b.vanwerkhoven@esciencecenter.nl

For questions about Atos-Bull, please contact:

Dr. Erwan Raffin, senior HPC expert

Tel.: +33 4 76 29 81 14

Email: erwan.raffin@atos.net

4.2 Other information

A separate Review Committee (RC) will be set up for the evaluation of the proposals submitted in this call. The RC will consist of scientific and technological experts from both within and outside of the ESIWACE2 Consortium. Members of the RC, and colleagues within their research group, are not allowed to submit a proposal in this call.