

Requests for Support

A CALL FOR PROPOSALS BY THE NATIONAL ESM SUPPORT TEAM

2022, VERSION 3



SPONSORED BY THE



Contents

1. Introduction.....	2
1.1 Background to the natESM initiative.....	2
1.2 Purpose of the natESM Support Team	2
1.3 Scope of the support services	3
2. Guidelines for applicants.....	3
2.1. Who can apply?	3
2.2. What can be requested?	4
2.3. When can requests be submitted?	4
2.4. Preparing a request for support.....	4
2.5. Submitting a request	4
2.6. During and after implementation of a support service.....	5
3. Assessment procedure	5
3.1. Procedure	5
3.2. Admissibility and evaluation criteria	6
4. General information	7
5. Contact	8

1. Introduction

1.1 Background to the natESM initiative

Global environmental change requires a gradual improvement in our ability to simulate, understand and predict the Earth system and its impact on society. Earth system modelling (ESM) can benefit enormously from access to exascale high-performance computing (HPC) systems, which are expected to become available in Europe and Germany in the mid-2020s. However, the use of the new HPC systems also poses significant challenges for model development in general and software development in particular. Since the new systems are likely to be based on heterogeneous hardware architectures and will use accelerators together with or instead of conventional CPUs, it is often not possible to continue using the code base of the models without fundamental adaptations.

During several workshops of the German ESM community that have taken place since 2018, a consensus was reached that the resources in the home institutes of the ESM community are not sufficient to prepare such fundamental code changes. Of course, criteria have to be met in order to develop a next-generation ESM system that allows progress in HPC to be used in the long term.

Therefore, a new level of ongoing inter-institutional collaboration is required that pools national resources in such a way that scientific and technological developments can be accelerated and shared. The national ESM strategy aims to achieve this collaboration. As a first step, the Federal Ministry of Education and Research (BMBF) has therefore funded the "natESM support team" project, the main aim of which is to address the aforementioned steps. In addition to four positions for research software engineers (RSE) at DKRZ and JSC, the position of coordinator will also ensure the targeted work of the national ESM strategy. Initially, this offer is funded for the period 2021-2025, but the aim is to establish a permanent institution.

1.2 Purpose of the natESM Support Team

HPC is the workhorse for the ESM community, providing new insights into climate processes to better understand climate change and assess adaptation scenarios. Although climate models are becoming more accurate and trustworthy, their current scope is often still limited. On the one hand, this is due to the missing or inaccurate representation of small-scale processes. On the other hand, however, it is also due to the fact that current models do not yet fully exploit the possibilities of HPC. In addition, the HPC landscape has changed dramatically in recent years, making code adaptations even more difficult and becoming an ongoing task. Programmers and even users have to deal with different top-level systems (CPUs, GPUs, vectors) and a variety of nuances in the details of usability and programming (x86 vs. ARM, OpenACC vs. OpenMP, etc.). The code must be flexible to take advantage of heterogeneous HPC systems and must be transferable between different sites - from small institutional clusters to medium-sized national centres to large European tier-0 systems.

All these issues motivate to jointly address the challenges on new HPC systems and through code refactoring in order to enable the necessary steps towards a national German ESM strategy. To reduce the effort of adapting or rewriting each individual code, a common infrastructure (programming paradigms, software libraries and finally the entire code base) would be desirable. There is a clear need to establish common ways of thinking and working (at least in terms of best practices) to prepare for exascale and to define common interfaces for coupling the different ESM components. For this reason, the project was funded and offers targeted support for the adaptation of Earth system models via the support team, as such comprehensive adaptations certainly cannot and should not be carried out repeatedly in a single institution.

While the main scientific development and optimisation work in the ESMs still needs to be done by the "home" institutions of the component models, the more technically oriented issues around HPC and software development can be addressed jointly. Appropriately embedding the components in a national strategy also requires serious efforts beyond institutional strategies. Incentives should be created by demonstrating that the new strategy provides resources for technical and organisational infrastructure, encourages cross-institutional collaboration, and offers support and training for groups wishing to transition their models.

1.3 Scope of the support services

The establishment of a national ESM support team will help to save resources, awaken synergies, share insights and disseminate knowledge. We aim to

- deliver direct HPC-support and consultancy to the ESM community,
- build up an expertise in the focus area: HPC, workflows, modularisation/scalable coupling of components, scalable diagnostics,
- offer training program for the ESM community to prepare codes independently in the home institutes.

The current funding phase of the national ESM support team must clearly be seen as a preparatory phase to prepare codes for new HPC systems and programming paradigms. The support provided by the support team in the work on Earth system models essentially comprises the necessary steps to enable subsequent upscaling HPC use. Initially, all steps in the software development cycle can be tackled jointly with the main developers of the models, e.g.

- Analysis and profiling of existing codes,
- Proof of concept implementation of new approaches,
- Porting of (parts of) existing code to new HPC architectures.

In the further course of the national ESM strategy, it is envisaged that the support work of the support team will also include topics that require the adaptation or integration of the individual models into an overall concept of national Earth system modelling. This can be done, among other things, through the following steps

- Definition of interfaces between model components,
- Use of uniform infrastructures (e.g. coupler, I/O, build and run environments),
- Setting up test environments.

2. Guidelines for applicants

2.1. Who can apply?

This call for proposals is open to all German model development groups in the field of Earth system research, including atmospheric sciences, oceanography and climate-related fields such as land or sea ice modelling, land surface modelling, atmospheric or ocean biogeochemistry, etc. Global models as well as regional models or small-scale, very high-resolution applications are addressed.

Since the funding is provided by the Federal Ministry of Education and Research (BMBF) and is intended to support the establishment of a national strategy, priority will be given to model codes whose main development is taking place at German research institutions.

We aim for a lasting improvement of the application's source code, and this is most likely to be achieved if the applicant is a member of the model code's core development team. Derived

"downstream" versions or user-specific branches have a lower chance of being approved, as we do not want to contribute to the fragmentation of the national Earth system modelling landscape.

2.2. What can be requested?

The support service covers up to 6 person-months per proposal (depending on the type and difficulty of the request) in the form of consulting and programming work by the RSE experts for high-performance and heterogeneous computing from DKRZ and/or JSC. The work is either performed by a colleague on a dedicated basis or as part of the support team and is basically to be understood as "help for self-help". The team can only help in an appropriate way if a complete task is not handed over for sole work, but if there is always an appropriate contribution from the requesting institution. Although the engineers will work largely remotely, we expect the applicant's research group to make appropriate efforts to ensure cooperation with sufficient support for the RSE. Similarly, there must be a minimum of prior work on current HPC systems such as scalable MPI parallelisation to address large distributed memory architectures.

With regard to the scope of the requested support, we basically distinguish between two categories

1. Short advisory activities or preparatory enquiries lasting less than 2 months. These include, for example, an initial status assessment of the code (profiling and analysis of shortcomings with regard to upcoming HPC systems) and can be seen as preparatory work for a later more comprehensive application.
2. Detailed assignment of special tasks with a term of up to 6 months. This includes, for example, the preparatory programming work necessary to run a model on GPU systems, or the implementation of conceptual programming paradigms such as hybrid parallelisation or modularisation.

A mutual agreement on the scope of the work to be done by the support team is reached in advance and the necessary time is measured on this basis. If the work is not completed within the planned time frame, it must be decided on a case-by-case basis whether an extension is possible. In principle, however, a follow-up application should be considered so that other requests can also be processed.

2.3. When can requests be submitted?

Applications can be submitted at any time and will be processed according to the procedure in [chapter 3](#). Should it become apparent during the project that a dedicated application phase (e.g. only twice a year) would be more helpful, we reserve the right to adjust the procedure. Information on this will be published on the website www.nat-esm.de/services.

2.4. Preparing a request for support

An application can be submitted via the template – published on the website <https://www.nat-esm.de/services/request-for-sprints>.

Applications must be written in English. As indicated on the webpage, the audience for this document is the ESM steering committee and technical support team.

Access to the model code and, if applicable, the documentation must already be available for the technical assessment of the application.

2.5. Submitting a request

In the current phase of the project, we assume that a submission of the documents (in the form of a single pdf document) by e-mail is sufficient: support-request@nat-esm.de.

If necessary, the process can be adapted in the further course of the project and will be announced via the website www.nat-esm.de/services.

In case certain points cannot be addressed, it is advisable to contact the natESM support team in advance. This exchange before submitting a request can facilitate the process for both sides. If you have specific questions about the request process or the selection process, please contact: info@nat-esm.de.

2.6. During and after implementation of a support service

The applicant is expected to have a basic infrastructure for remote software development ready or to establish it in the first weeks of the cooperation. This includes a version control system, as well as a platform for tracking issues, progress and discussions. If this cannot be provided by the applicant itself, it should be checked whether, for example, DKRZ's gitlab.dkrz.de service is a solution. Likewise, representative test cases are necessary to monitor the numerical results and performance of the model.

As a significant part of the collaboration will only take place virtually, the applicant is expected to show an open and collaborative way of working during the development cycle. This means, among other things, responding promptly to emails and giving advance notice of absences. When RSEs visit the applicant's institute, the applicant is responsible for providing the RSEs with a suitable, safe and healthy working environment.

If these conditions are systematically violated, the steering committee of the natESM initiative can decide to terminate the support service at any time. In particular, the effectiveness of the cooperation should be assessed in the first month of the cooperation and the project should be terminated if it is not effective.

The applicant (with support team colleagues and natESM coordinator) writes a short report describing the progress made through the collaboration, the impact on the performance or transferability of the code and the impact on the scientific use cases. This report is due within 3 months after the end of the support period. There is a template for this report, which highlights what successes or failures have occurred. It will be posted on www.nat-esm.de/services/documentation after submitting.

3. Assessment procedure

3.1. Procedure

The entire process runs in parallel in two stages (technical and scientific comments according to Figure 1 and [section 3.2](#)) and should normally be completed within 4 weeks. The technical judgement of the proposal will be carried out by the RSEs team, although other colleagues from DKRZ and/or JSC may be called in if necessary. It will not assess the scientific impact, but rather the technical feasibility of the project. In case of ambiguities and follow-up questions during the technical assessment process, the applicant is invited to a video conference.

The accompanying scientific comments include in particular the relevance of the model, its consistency with the national ESM strategy and the assessment of whether the scientific issue presented can be adequately solved with HPC systems. Depending on the sub-discipline, members from the Steering Group will be active and/or additional renowned experts from the fields of Earth system research with affinity to HPC will be consulted. Queries from the reviewers are possible and should be answered promptly via mail, phone or, if necessary, video conference. The final decision on the application will be made by the Steering Group of the National ESM Initiative.

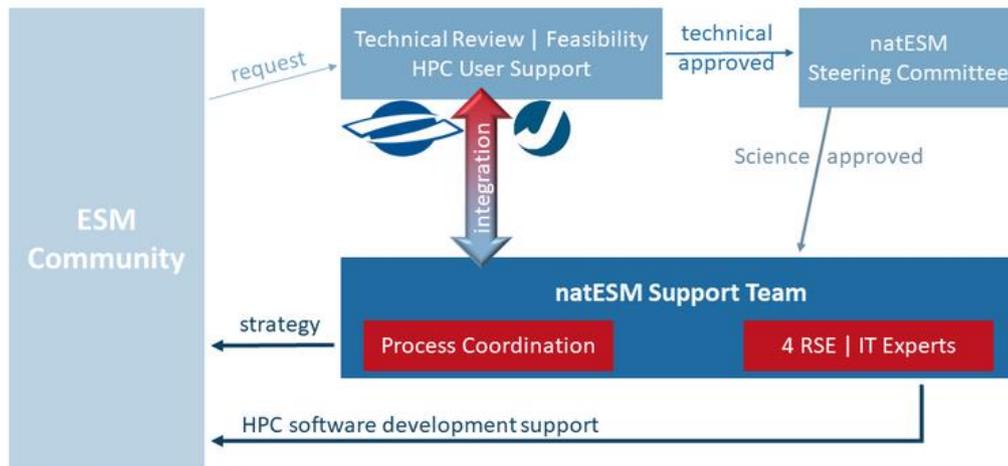


Figure 1: Procedure of the assessment process for support requests

3.2. Admissibility and evaluation criteria

A proposal is considered eligible if the applicant belongs to a German non-commercial research institution that meets the description under 2. Furthermore, the code must be freely accessible in the assessment process and later on. Finally, the proposed hardware platform must exist in Europe and access to the hardware for both the applicant and the natESM Support Team should be plausible.

It must be made clear that the necessary computing time to carry out the support service is not part of this project and must be provided by the applicant.

The technical comments will be based on the following criteria. This list is not comprehensive and can be adapted if necessary. If certain points are not fulfilled or not completely fulfilled, this is not a criterion for exclusion, as technical further developments in particular can be part of the support work.

1. Formulation of realistic goals
 - Is the targeted (performance) improvement realistic, can the code e.g. be transferred to a new architecture within the proposed timeframe?
 - Can the underlying numerical algorithm and/or code implementation be accelerated at all?
2. Software perspective and impact
 - Does the proposal provide a credible path for long-term goals of the software?
 - Will the proposed changes be incorporated into the main branch of the model and will this code be sufficiently maintained afterwards?
3. Formal technical requirements
 - Is there useful, user-friendly documentation of the code?
 - Are standard methods (like programming paradigms) followed?
 - Is there a licensing model that allows open source distribution?
 - Is there already a system for traceability, reproducibility and version control?
 - Are well-defined interfaces between the components of the Earth System provided?
 - To what extent is the code already scalable, portable, and has aspects of modularity?

The scientific comments will be based on the following criteria. This list is not comprehensive and can be adapted if necessary. If certain points are not fulfilled or not completely fulfilled, this is not a criterion for exclusion.

1. Quality and scale
 - Are simulations from global to local made possible?
 - Is there a scaling workflow?
 - Is there capacity for data assimilation and diagnosis?
2. Improvements and innovations
 - Does the proposed work improve the quality of simulations and enable its users to achieve unprecedented accuracy?
 - Do the adaptations open up new scientific areas or do they hold the potential to go beyond the current state of the art?

4. General information

As the work is done on a project basis, no claim to long-term maintenance of the contributed code can be made from a support grant once awarded. Long-term maintenance will be undertaken by the applicant's modelling group. The contributed code will be written and documented according to the rules of good scientific practice or the internal coding standards of the modelling group.

Any results or concepts within the scope of the support service can be disseminated by the natESM support team at conferences and workshops as well as reused for further work within the project. On the one hand, this ensures the long-term re-usability of the results, but on the other hand, it also offers the modelling group the possibility to disseminate their model. If legal steps stand in the way of this procedure, this must already be explained when applying for the support service and will have an impact on the selection process.

The natESM support team believes that science must be as open as possible. Results from this publicly funded project must be published Open Access, all relevant data must be made available as Open Data and software as Open Source. Exceptions may arise, for example, in the case of a data set that is sensitive under data protection law or the necessary use of protected software. In principle, we have a strong preference for developing code as open source with a permissive licence.

We do not claim intellectual property rights (IPR) to the jointly developed code. Intellectual property rights to the code remain with the applicant institution, but the natESM support team reserves the right to use the results in other/further support services. This is to implement the idea of the collaborative natESM strategy.

The collaborative partners of the support service shall carry out the agreed work with their usual care, based on the state of the art in science and technology known to them and in compliance with the guidelines for ensuring good scientific practice.

Claims of the partners against each other for compensation for damages arising from breaches of duty and from tort are excluded unless they are based on intent or gross negligence. Liability for indirect and consequential damages, in particular loss of profit, is excluded.

5. Contact

If you have specific questions about the support request process, the selection process or the natESM project in general, please contact:

info@nat-esm.de

For questions about the partnering computing centres involved in the natESM project, please contact:

Dr. Hendryk Bockelmann (for DKRZ)

Tel.: +49 40 460094-144

Email: bockelmann@dkrz.de

Dr. Sabine Griessbach (for JSC)

Tel.: +49 2461 61-6555

Email: s.griessbach@fz-juelich.de