



Project scope and plan

Project name	please fill in your project name
Research field	please fill in your research field

Principal Investigator (PI)

Title (Dr., Prof., etc.)	Dr., Prof., etc
Last name	please fill in your last name
First name	please fill in your first name
Organisation name	e.g. University of Amsterdam
Department	e.g. Department of Physics
Group	e.g. Molecular Physics Group
Country	please fill in your country

Co-PIs (same information)

Provide the details of any Co-PIs in the project, including Title, Last name, First name, Organisation, Department, Group and Country.

It is mandatory to include all team members in the online form.

Co-Principal Investigator (Co-PI)

Title (Dr., Prof., etc.)	Dr., Prof., etc
Last name	please fill in CoPI last name
First name	please fill in CoPI first name
Organisation name	e.g. University of Amsterdam
Department	e.g. Department of Physics
Group	e.g. Molecular Physics Group
Country	please fill in CoPI country

IMPORTANT NOTICE

All of the sections and subsections below **MUST BE COMPLETED** (unless stated otherwise). In case you wish to leave a section empty, please provide a reason. The Access Committee (AC) will not be able to process proposals that neither provide the requested information nor a justification for the lack of such information for each section.

Applicants are strongly encouraged to **base their proposal on reliable benchmark data obtained on the target machine(s)** from previous calls order access programmes. Such data and support to properly collect these can be obtained from the PRACE Preparatory Access. If needed, please contact the peer review office of PRACE at peer-review@prace-ri.eu in order to request such access and support, **at least 1 month before the submission deadline**.

The structure and formatting settings of this template must be preserved and respected (change in font size or margin and spacing settings are not allowed). The maximum number of pages allowed is 20 pages, including graphs, tables and references, but not counting the cover page and the appendix. Reviewers will be instructed not to consider any pages out of the limit. Applicants are requested to include information about their track record in Appendix 1 at the end of the document (not counted in the page limit). **Instruction paragraphs can be removed from the proposal text by disabling the 'hints' option to the class.**

Upload a single document, based on the present template, in PDF format **without exceeding** 8 MB.

Proposals that do not follow the template or that are incomplete will be administratively rejected and will not be further evaluated.

1 Key scientific/societal/technological contribution of the proposal (200 words max.)

Outline the scientific/societal/technological importance of your project, how High Performance Computing (HPC) will help you achieve your goals and what the major expected outcomes are. This section would typically be the same as the abstract of the proposal in the submission form.

2 Detailed proposal information (maximum 14 pages, graphs and tables included)

The information should be suitable for expert peer review in your discipline. It must also have appropriate information for a broader audience as your proposal will be evaluated by a panel and in parallel with proposals in other disciplines.

2.1 Justification for the importance of the scientific problem and the requested resources (~2 pages)

Describe the proposed research and the main scientific/technical advances you will achieve with the requested PRACE allocation. Industry organisations should also summarize the potential economic or strategic business impact. The justification of the requested resources must be clearly linked to the software performance evaluation (Section 2.6).

2.2 Overview of the project (~4 pages)

Describe the motivation, objectives and scientific challenges of the problem. Describe and justify the choice of computational methods. State the advances that will be enabled through the requested Tier-0 PRACE award (e.g.

impact on community paradigms, valuable insights or solving a long-standing challenge, new technology/therapy, etc.). Provide a list of expected outcomes of your proposal and, if relevant, the interdisciplinary value of your proposal.

2.3 Validation, verification, state of the art (~2 pages)

Please describe the validity of the simulations and predictions made with this proposal. In case you provide references to relevant publications please include here also the key relevant results. Please address issues of reproducibility and highlight the predictive capabilities of your simulations.

2.3.1 Validation

Validate your model against experiments or other established reference data.

2.3.2 Verification

Verify the numerical consistency of your method or provide evidence of existing verifications.

2.3.3 Sensitivity analysis and uncertainty quantification

Provide sensitivity analysis of your method. Provide estimates of the uncertainty of your predictions. Data driven uncertainty quantification is encouraged. In case of multiphysics/multiscale problems, uncertainty of the full methods and software is desirable.

2.3.4 Comparison with state of the art

Place the project in the context of competing work. Explain the relative advantages AND drawbacks of your approach.

2.4 Software and Attributes (~2 pages)

(Please see also **Examples of Performance Reporting** in Section 2.6.2.1). Describe the software that will be used including a discussion of the state of the art in the field. The description should mention:

2.4.1 Software

Please describe all codes you are using in the proposal. Justify your choices and describe alternatives (if any).

2.4.2 Particular libraries

Describe particular libraries required by the production and analysis software, algorithms and numerical techniques employed (e.g. finite element, iterative solver), programming languages. Please specify requirements for compilation or build environment (build system (e.g. cmake, python version), version control system (e.g. git, subversion) etc.).

2.4.3 Parallel programming

Model(s) used (e.g. MPI, OpenMP/Pthreads, CUDA, OpenACC, etc.).

2.4.4 I/O requirements

I/O requirements (e.g. amount, size, bandwidth, etc.) for execution, input files, restart and other output. Describe I/O strategy (number of files, frequency, read/write size) and I/O behaviour of your code during the period of calculations. Please specify the restart overhead, not only for I/O; (e.g. a code may have to perform a costly domain decomposition first).

2.5 Data: Management Plan, Storage, Analysis and Visualisation (~1 page)

2.5.1 Data Management Plan covering

Data Management Plan covering both short-term and long-term aspects, including needs for I/O bandwidth, number of files and input/output data volumes. Specify for which system the data will be provided, how long the data must be stored at the computing centre after the termination of the project, how it will be moved from the centre, and how subsequent analysis will be performed. Specify the availability of both code and data to other researchers, and how this will be handled. PRACE should be given credit for all data produced through PRACE allocations when publishing, and described in the provenance when depositing to other infrastructures.

2.5.2 Project workflow

Project workflow including the role and timeline of data analysis and visualisation identify where the analysis will be done and any potential bottlenecks in the analysis process. Describe any analysis and/or data reduction tools used.

2.5.3 Software workflow solution

Software workflow solution (e.g. pre- and post-processing scripts that automate run management and analysis) to facilitate this volume of work.

2.5.4 I/O requirements

I/O requirements (e.g. amount, size, bandwidth, etc.) for data analysis and visualisation. Highlight any exceptional I/O needs. Please provide data for (one or several) precise systems that will be simulated.

2.6 Performance of Software (Maximum 3 pages)

2.6.1 Testing of your code on the requested machine

It is strongly recommended that your production code is tested in the requested machine(s) (see also the text referring to Preparatory Access in the Important Notice at the top of Page 1). Please specify the PRACE Preparatory Access project (if any) or other projects (previous PRACE Calls, national calls, etc.) used to prepare the Tier-0

proposal. If the preparatory host machine is different from the target machine, specify why you think the data is relevant. In the latter case, please report briefly the conversion factor (in terms of ratio of time to solution, flops or requested core hours) from the machine where the preparatory tests were performed to the requested system. Moreover your proposal must account for all technical constraints and requirements of the targeted machine(s) as documented in the separate *Technical Guidelines for Applicants* document; failing to do so will significantly increase the risk that your project will be technically rejected.

2.6.2 Quantify the HPC performance of your project

The presented data must be representative of the entire workflow of the project proposed and refer to the main application code you intend for the production work. The software scalability data (see **Examples of Performance Reporting** below) must be used to choose the most efficient job size(s) for the simulations planned in production: the corresponding software performance must be clearly linked to the justification of the computing resources requested. The AC will not accept estimates based on related codes and/or data related to parts of your production. All data must refer to the targeted systems in your production runs or a system with comparable size, software stack and with the same architecture, and network (e.g. a project can be accepted on Marconi-KNL if it was benchmarked on another KNL machine with the same Omnipath network). Please coordinate with the centres if in doubt about the portability of your code. Specify that performance results are reported on the basis of one of the following: whole application including I/O; whole application except I/O; kernel only; other (specify). More specifically you must include:

2.6.2.1 Strong and weak scalability

Starting with the minimum size of the computer necessary to run the problem (usually 1 core or 1 node). Justify the minimum size for your scaling if it is larger than 1 core or 1 node (e.g. memory limitations). Please provide a justification in case that either the weak (e.g. study of one particular bio-molecule) or the strong (e.g. study of an ensemble) scalability metric is not considered relevant to your project. See **Examples of Performance Reporting** below for the requested format.

Examples of Performance Reporting. For the weak and strong scaling please start with the minimum and finish with the maximum number of cores that are suitable for your application. Please mark the number of cores that you expect to perform the main load of your work. On the Y axis you may use time to solution (scaled or otherwise) or speedup with respect to the minimum number of cores. *The table with the timings is mandatory.* The table should include the speedup and the parallel efficiency. Log/log plots are useful to span many orders of magnitude.




Table 1: Weak scaling example data

Number of cores	Time to solution (s)	Ideal time to solution (s)
4,800	10.50	10.50
9,600	10.55	10.50
19,200	10.58	10.50
38,400	10.60	10.50
76,800	10.57	10.50
153,600	10.57	10.50

Figure 1: Weak scaling example figure




Table 2: Strong scaling example data

Number of cores	Time to solution (s)	Ideal time to solution (s)
4,800	5,000.00	5,000.00
9,600	2,725.00	2,500.00
19,200	1,500.00	1,250.00
38,400	1,000.00	625.00
76,800	700.00	312.50
153,600	500.00	156.25

Figure 2: Strong scaling example figure

2.6.2.2 Precision reported

One of: single precision, double precision, mixed precision. Only the precision you use in the simulation is relevant.

2.6.2.3 Time-to-solution

Time-to-solution as normalized/averaged per iteration,

$$T_i^* = \frac{t_i \cdot N_c}{N_e} \quad (1)$$

AND normalized as total time to solution,

$$T_f^* = \frac{t_f \cdot N_c}{N_e} \quad (2)$$

with t_i the time per iteration, t_f the total time to solution, N_c the number of cores and N_e the number of computational elements (size of the problem).

IMPORTANT: Justify the choice of your code (e.g. comparison with existing codes, methods or any other scientifically rigorous argumentation). See also the text referring to Preparatory Access in the Important Notice at the top of Page 1.

2.6.2.4 System scale

One of: results measured on full-scale system, projected from results of smaller system, other (specify).

2.6.2.5 Measurement mechanism

One of: timers, FLOP count, static analysis tool, performance modelling, other (specify).

2.6.2.6 Memory usage

Specify requirements per node or core depending on the size of the problem.

2.6.2.7 OPTIONAL: Percentage of available peak performance

Please collaborate with the Centres on obtaining this information (see also the text referring to Preparatory Access in the Important Notice at the top of Page 1). Alternatively provide code specific metrics for the requested machine (FLOPS, etc.).

3 Milestones (quarterly basis) (maximum 1 page)

Goals and milestones should articulate simulation and developmental objectives and be sufficiently detailed to assess the progress of the project for each year of any allocation granted. It is especially important that you provide clear connections between the project's overarching milestones, the planned production simulations, and the compute time expected to be required for these simulations. Please clarify any dependencies of milestones on other milestones. Please ensure that the core hour consumption is regular throughout the allocation or provide a requested schedule after consultation with the centres.

3.1 Gantt Chart

Provide a Gantt Chart of the simulation plan in production indicating job sizes and scheduling of computing tasks.

3.2 Communication plan

Provide a communication plan for the results, describing the strategy and timeline for scientific publication, conference presentation (e.g. given by the PI at PRACEdays), and other public outreach events.



4 Personnel and Management Plan (~1/2 page)

What personnel are already in place and what are their roles on the project? If applicable, describe (i) personnel that will be hired for the project in the future and their responsibilities and (ii) potential personnel turnover that may occur during the project and a strategy for replacing them. The PRACE program welcomes proposals from individual PIs or teams of collaborators. Outline the focus of each individual or subgroup and their interrelationships.

It is mandatory to include all team members on the online form.

5 Previous Allocations and Results (~1/2 page)

Please specify previous PRACE allocations and relevant results. Any results of relevance to the project need to be listed to demonstrate how the proposal contributes to the long-term goals of the proposer. Research publications and reports that resulted from previous allocations should be listed in the References below. Please specify in particular references to publications that acknowledge the use of PRACE resources.

6 References (maximum 30)



Appendix 1: Track Record of the PI

Curriculum vitae and list of publications of the Principal Investigator

Granted patents and other measures for the relevance of the work

Prior allocation history in PRACE, national calls, as well as international programs such as INCITE of the US DoE

Participation by team members in other European Commission (EC) actions, such as ERC or Marie Skłodowska-Curie EC grants, etc.

Previous presentations at PRACEdays
