PRACE Project Access
Peer Review Process and Proposal Template

PRACE Dissemination

Innsbruck, January 21st 2019

Maria Grazia Giuffreda and Luca Marsella (ETHZ-CSCS)

Outline

The PRACE Peer Review of proposals for a PRACE Tier-0 systems combines **scientific** and **technical** assessments

The assessments are carried out by two groups of experts: the **scientific** review **builds** on the **technical** one

We will discuss the key elements of a successful proposal, considering both the scientific and the technical aspects

We will illustrate the **proposal template**, highlighting the **technical** information **necessary** for a successful proposal

PRACE provides a federate European supercomputing infrastructure that enables **large computational projects** in science

A strict and strong transparent **peer review process** based on **scientific excellence and HPC competence** is in place to allocate resources on the world-class Tier-0 supercomputing infrastructure

Two PRACE Tier-0 Calls per year

Provisional Planning

Call 18 (Closed)

- Opening of the call: 4 September 2018
- Closing of the call: 30 October 2018, 10:00 CET
- Allocation period: From 2 April 2019 to 1 April 2020

Call 19

- Opening of the call: 5 March 2019
- Closing of the call: 30 April 2019, 10:00 CEST
- Allocation period: From 1 October 2019 to 30 September 2020

Project Access

Proposal duration: 12-months schedule (**Single-year Projects)**

24- or 36-months schedule (**Multi-year Projects)**

Even multi-year projects are initially granted for 1 year only. Allocations for 2nd and 3rd year are tentative

Centres of Excellence (CoE):

0.5% of all resources are reserved for CoEs

CoEs are selected by the European Commission under the E-INFRA-5-2015 call for proposals

Who can apply?

Scientists and researchers from academia and industry:

The project leader must have an employment contract as researcher in the organization at the time of proposal submission

Only research proposals of a civilian (non-military) nature are eligible

Double-granting is not permitted:
Proposals admitted to any other HPC program will be rejected

Resources: what to apply for

System	Architecture	Site (Country)
Joliot Curie SKL	Bull Sequana X1000	**GENCI@CEA (FR)**
Joliot Curie KNL	BULL Sequana X1000	**GENCI@CEA (FR)**
Hazel Hen	Cray XC40	**GCS@HLRS (DE)**
JUWELS	Multicore Cluster	**GCS@JCS (DE)**
SuperMUC	Lenovo NextScale	**GCS@LRZ (DE)**
SuperMUC-NG	Lenovo ThinkSystem	**GCS@LRZ (DE)**
Marconi-Broadwell	Lenovo System	**CINECA (IT)**
Marconi-KNL	Lenovo System	**CINECA (IT)**
MareNostrum	Lenovo System	**BSC (ES)**
Piz Daint	Cray XC50	**CSCS (CH)**

Peer Review Process (ERC Process - like)

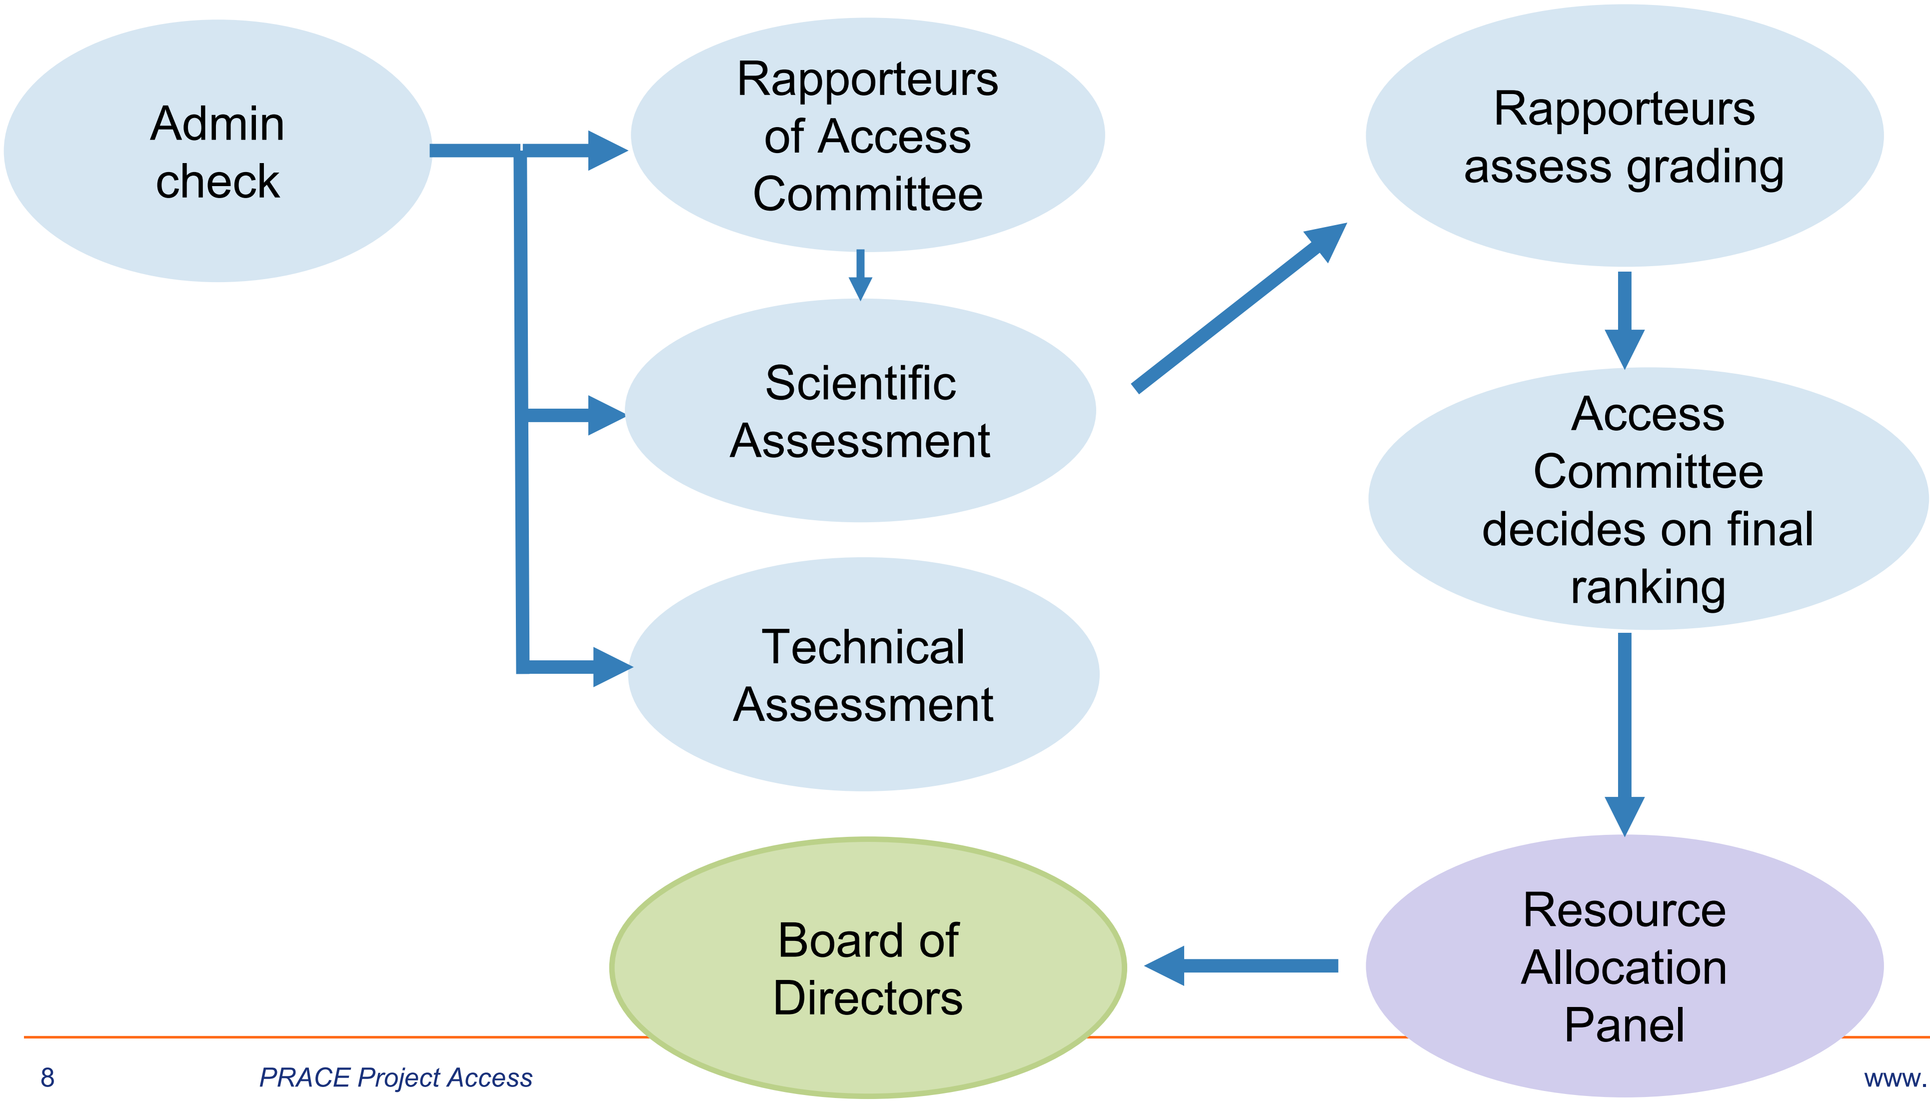

Peer Review Process

Administrative check

Proposals not complying with PRACE eligibility criteria will be rejected (e.g. CV or publication list missing, wrong template, too long, etc.)

Peer Review Process

Technical Assessment

Proposals will be reviewed by technical experts of PRACE Hosting Member sites to assess suitability

Applicants may be contacted by technical experts in case of questions or concerns raised during review

Technical data need to be provided for the targeted system

Peer Review Process

Scientific Assessment

Each proposal is assigned to three scientific peer reviewers to assess

- scientific merit
- soundness of numerical methods, algorithms, and computational tools
- appropriateness of project timeline and resources
- feasibility of research plan
- qualifications, expertise, and track record of PI and team

Peer Review Process

Access Committee (AC)

Each proposal is assigned to two members of the PRACE AC as Rapporteur 1(R1) and Rapporteur 2 (R2), who:

- suggest scientific peer reviewers and solicit reviews
- assess returned reviews
- discuss discrepancies in the grading of the proposals
- agree on common grade

Scientific excellence is the **single decisive criterion**

Peer Review Process

Access Committee (AC)

Proposals are ranked according to the grade of the R1 and R2 consensus

Proposals below a certain threshold are rejected unless any member of the AC wishes to discuss them further

Peer Review Process

Access Committee (AC)

Every proposal is discussed in the presence of the full AC as follows:

- R1 presents the proposals objectives, strengths and weaknesses, reports grade
- Technical members discuss the feasibility of the project
- R2 makes further remarks and may propose an alternate grade
- All AC members discuss the proposals and ask questions to R1 and R2
- R1 proposes a score followed by a score by the AC Chair
    - in case of agreement this become the tentative score of the proposal
    - in case of disagreement there is a vote

Peer Review Process

Access Committee (AC)

Proposals are re-discussed when similarities are found with other proposals and grades may be revised

Final AC session: discussion and provision of the final ranking of all proposals

Peer Review Process

The Resource Allocation Panel

This is the final step of the review process:

The panel decides on allocations based on

- recommendation of the AC
- constraints on PRACE resources
- administrative constraints and agreements of PRACE partners

The PRACE BoD confirms the final allocations

Scientific Assessment

- Scientific Excellence
- Novelty and transformative qualities
- Relevance to the call
- Methodology
- Dissemination
- Management

Scientific Assessment

- **Scientific Excellence**
  - The proposed research must demonstrate scientific excellence and high **European** and **International impact**

Scientific Assessment

- **Novelty and transformative qualities**
  - The proposed project should develop transformative topics of major relevance to **European** research

Scientific Assessment

- **Relevance to the call**
  - The proposal should describe how the research is addressing the **scope** of the **call** if a specific scope is stated in the call

Scientific Assessment

- Methodology
  - The **mathematical numerical** methodology should be described and be appropriate to achieve the project's goals

Scientific Assessment

- ## Dissemination
  - The **channels** and **resources** for dissemination should be described. The list of recent relevant publications is **essential**

Scientific Assessment

- Management
  - There must be a **solid management** structure in place, ensuring that the project will be successfully completed

Proposal Template

You should provide information on all of the subsections

- If you wish to leave a section empty, provide a reason!
- The information should be suitable for expert review in your field
- …but also appropriate for a broader audience: your proposal will be evaluated by a panel with proposals in other disciplines

Key contribution of the proposal

- **Scientific / Societal / Technological contribution**
  - Outline the societal importance of your project
  - How will HPC help you achieve your goals?
  - What are the major expected outcomes?

Importance of the scientific problem

- **Justify the scientific relevance and the resource request**
  - Describe the proposed research and the main scientific / technical advances you will achieve with the requested PRACE allocation
  - Industrial partners 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

Overview of the Project

- Describe the motivation, the objectives and the scientific challenges, justifying the choice of computational methods
- State the advances enabled through the requested Tier-0 PRACE award (e.g.: impact on community paradigms, new insights, etc.)
- Provide a list of expected outcomes of your proposal and, if relevant, the interdisciplinary value of your proposal

Validation, Verification, State of the art

- Describe **validity** of simulations and predictions resulting from resources
- Provide **references** to publications and address their **reproducibility**
  - **Validation**
  - **Verification**
  - **Sensitivity analysis and uncertainty quantification**
  - **Comparison with state of the art**

Validation, Verification, State of the art

- Describe **validity** of simulations and predictions resulting from resources
- Provide **references** to publications and address their **reproducibility**
  - **Validation:** Validate your model against experiments or other established reference data (if available)

Validation, Verification, State of the art

- Describe **validity** of simulations and predictions resulting from resources

- Provide **references** to publications and address their **reproducibility**

  - **Verification:** verify the numerical consistency of your method or provide evidence of existing verifications

Validation, Verification, State of the art

- Describe **validity** of simulations and predictions resulting from resources

- Provide **references** to publications and address their **reproducibility**

  - **Sensitivity analysis and uncertainty quantification:**
    - Provide **sensitivity analysis** of your methods
    - Provide estimates of the **uncertainty** of your predictions
      Data-driven uncertainty quantification is encouraged
    - In the case of multiphysics / multiscale problems, **uncertainty** of **methods** and **software** is desirable

Validation, Verification, State of the art

- Describe **validity** of simulations and predictions resulting from resources

- Provide **references** to publications and address their **reproducibility**

  - **Comparison with state of the art:**
    - Place project in the context of competing work
    - Explain advantages **AND** drawbacks of approach

Software and Attributes

- Describe the software that will be used including a discussion of the state of the art in the field
  - **Software**
  - **Particular libraries**
  - **Parallel programming models**
  - **I/O requirements**

Software and Attributes

- Describe the software that will be used including a discussion of the state of the art in the field

  - **Software**
    - All codes you are using in the proposal
    - Justify your choices and describe alternatives

Software and Attributes

- Describe the software that will be used including a discussion of the state of the art in the field

  - **Particular libraries**
    - required by production analysis software, algorithms and numerical techniques, programming languages

Software and Attributes

- Describe the software that will be used including a discussion of the state of the art in the field

  - **Parallel programming models**
    - MPI, OpenMP/Pthreads, CUDA, OpenACC, etc.

Software and Attributes

- Describe the software that will be used including a discussion of the state of the art in the field

  - **I/O requirements**
    - amount, size, bandwidth, input files, restart and other output
    - Describe I/O strategy (number of files, frequency, read/write size)
    - I/O behaviour of the code during execution
    - Specify the restart overhead (e.g. costly domain decomposition)

Data Storage, Analysis and Visualization

- **Project workflow:** including **role and timeline** of data analysis and visualization; identify where the analysis is done and potential bottlenecks
  - Describe any analysis and/or data reduction tools used
- **Software workflow solution:** pre- and post-processing scripts that automate run management and analysis to facilitate volume of work
- **I/O requirements:** amount, size, bandwidth, etc. for data analysis and visualisation
  - Highlight any exceptional I/O needs
  - provide data for (one or several) precise systems that will be simulated

Software Performance

- **Information on software performance is mandatory**
  - production code should be tested on the requested machine. Specify the **preparatory project** (if any) or projects used to prepare the Tier-0 proposal
  - If the preparatory host machine is different from the target machine, then you need to specify why you think that the data presented is relevant
  - Report briefly the conversion factor (e.g. ratio of time to solution, flops or requested core hours) from the preparatory-test machine

Software Performance

- **Quantify the HPC performance of your project**
  - Data must be representative of the entire workflow of the project proposed and it must refer to the **main application code** for the production work
  - Scalability must be used to set most efficient job size for planned simulations and performance must be linked to the request of the computing resources
  - No estimates based on related codes or on parts of codes will be accepted
  - All data must refer to the targeted systems in production runs or a system with comparable size, software stack and with the same architecture and network
    - contact the Supercomputing Centre if in doubt about the code portability

Performance Results

- Strong and weak scalability
- Precision reported
- Time-to-solution
- System scale
- Measurement mechanism
- Memory usage
- Optional

Performance Results

- **Strong and weak scalability**
  - start with the minimum size of the computer necessary to run the problem **justify** the minimum size for your scaling if it is larger than 1 core or 1 node (e.g. due to memory limitations)
  - **justify** if **weak** or **strong** scaling is not relevant for the project (e.g.: weak scaling not relevant for the study of a particular biomolecule, strong scaling not relevant for an ensemble)

Performance Results

- **Precision reported**
  - E.g.: single precision, mixed
  - only the precision you use in the simulation is relevant

Performance Results

- **Time to solution**
  - normalized / averaged per iteration, number of cores and size of the problem: T_i^* = (Time-per-iteration) x (No. of cores) / (No. of computational elements)
  - normalized as total time to solution, number of cores and size of the problem: T_f^* = (Total-time-to-solution) x (No. of cores) / (No. of computational elements)
  - Justify the choice of your code (e.g. comparison with existing codes, methods or any other scientifically rigorous argumentation)

Performance Results

- **System scale**
  - measured on full-scale system
  - projected from results of smaller system
  - specify if other

Performance Results

- **Measurement mechanism**
  - timers, FLOP count, static analysis tool
  - performance modeling
  - specify if other

Performance Results

- **Memory usage**
  - specify requirements per node or core, depending on the size of the computational problem

Performance Results

- Optional
  - Percentage of available peak performance: collaborate with the Supercomputing Centre to obtain this information
  - Contact the peer review office of PRACE to request help of high-level support team (at least 1 month before deadline)
  - Alternatively provide code specific metrics for the requested machine (FLOPS, etc.)

Examples of Performance Reporting

- Start the scaling plots with the minimum simulation size and finish with the maximum number of cores suitable for your application
- Mark number of cores expected to perform simulations. On the Y axis use time to solution (scaled) or speed-up vs. 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

Examples of Performance Reporting

Weak Scaling Example

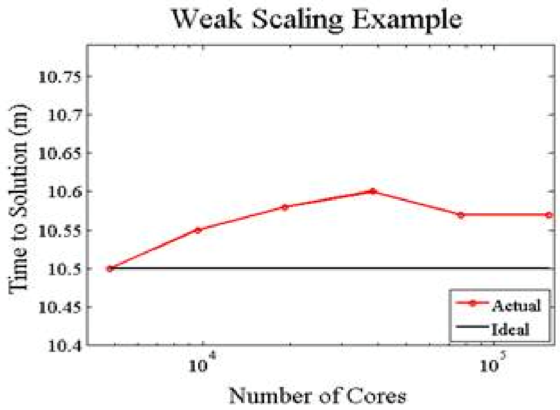

nProc	Time to Solution (m)	Ideal Time to Solution (m)
4800	10.50	10.50
9600	10.55	10.50
19200	10.58	10.50
38400	10.60	10.50
76800	10.57	10.50
153600	10.57	10.50

Strong Scaling Example

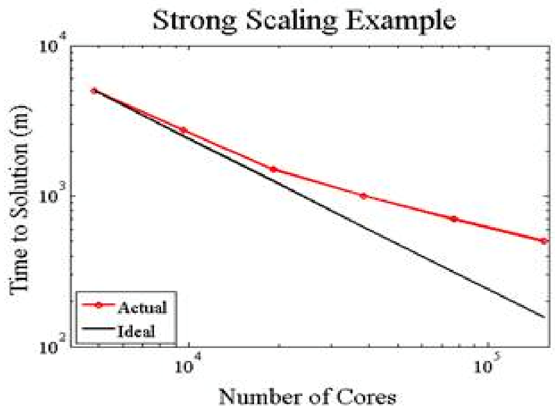

nProc	Time to Solution (m)	Ideal Time to Solution (m)
4800	5000.00	5000.00
9600	2725.00	2500.00
19200	1500.00	1250.00
38400	1000.00	625.00
76800	700.00	312.50
153600	500.00	156.25

Milestones (quarterly basis)

- Goals and milestones should articulate simulation and development objectives and be detailed to assess the progress for each year of allocation
- Important to provide clear connections between the project's milestones, the planned simulations, and the compute time required for such simulations
- Clarify any dependencies of milestones on other milestones
- Ensure that the core hour consumption is regular throughout the allocation or provide a requested schedule after consultation with the Computing Centres
    - Provide a Gantt Chart of the simulation plan in production
    - Indicate job sizes and scheduling of computing tasks in the chart

Personnel and Management Plan

Present personnel overview

- personnel that will be hired for the project and their responsibilities
- potential personnel turnover during the project and replacement strategy
- Specify if the proposal is from individual PIs or teams of collaborators
- Outline the focus of each individual or subgroup and their relationships

Previous Allocations and Results

- Provide references to publications that **acknowledge PRACE resources**

Further Information

- PRACE Application procedure
  - http://www.prace-ri.eu/application-procedure
- Available Tier-0 Supercomputing Systems
  - http://www.prace-ri.eu/prace-resources
- Project Scope and Plan
  - http://www.prace-ri.eu/IMG/docx/ProjectScopePlan_V6.docx
- Technical Guidelines for Applicants (updated for all 18)
  - http://www.prace-ri.eu/IMG/pdf/PRACE_Call_18_Technical-Guidelines_for_Applicants.pdf

THANK YOU FOR YOUR KIND ATTENTION

www.prace-ri.eu