PRACE 2.0 PROPOSAL for COUNCIL

VISION of PRACE

- To provide a world-class HPC infrastructure to European researchers and industry comprised at any one time of leadership-class pan-European systems, linked with an underpinning network of national and regional systems.
- Access to PRACE leadership-class systems is granted to all European researchers on the basis of a fair, equal and transparent peer review process based on scientific excellence and a fair cost sharing principle.
- Leadership-class systems need to be procured openly in a global market, following clear and transparent principles of user needs, technological leadership, and cost effectiveness.
- PRACE should be seen as the organisation providing leadership on HPC in Europe

PRACE 2.0

The aim of PRACE 2.0 is to continue providing all European scientific communities with adequate, leadership-class computing systems, whose resources are pooled into an integrated distributed forefront infrastructure, accessible via a centralized peer-review process. The specification of leadership-class systems is based on the requirements of different scientific disciplines. It is essential that PRACE 2.0 is built on the important achievements reached by PRACE 1.0, and that all members sign up to an explicit long-term goal for PRACE, so that a route to get there by 2020 can be developed.

The key challenge is putting in place sustainable funding to allow the persistence of European HPC infrastructure. It has to be recognised by funders that the renewal cycle and pattern of investment is different for HPC compared to other research infrastructures: it is NOT an one-time investment, but rather a continuous one.

A fair contribution of all partners in PRACE is required to build a sustaining PRACE model: shared participation of Hosting Members, the non Hosting Members and the European Commission in the operational costs of the infrastructure. The success of PRACE2.0 depends on the inclusion of all current partners and therefore PRACE aims at finding a cost model that is affordable for all members.
Hosting Members finance the required hardware investments nationally. The future model of contribution for OPEX is based on GDP and past usage of each member. A contribution from the EC is needed to make the model viable.

Scientific Case and Requirement

In 2015 PRACE will be offering a total of 18 Petaflops of peak performance on complementary architectures of 6 Tier0 systems located in France, Germany, Italy and Spain. This is similar to other RI for civil research in the US, Japan and China, allowing Europe to compete at the highest level within the HPC ecosystem.

Since 2010, PRACE has awarded 9.2 thousand million core hours to 346 projects (over a total number of 858 proposals leading to an oversubscription ratio of 248%) based on open peer reviewed calls for proposals with one single criteria: scientific excellence. This represents an average allocation of 26.5 million core hours, which is beyond what is usually available at the national level.

These 346 projects are covering all the scientific fields as shown in the following chart:

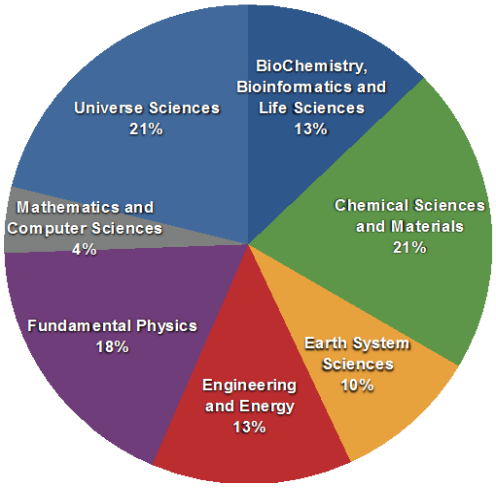

Many projects include international collaborations with the US, Japan, Russia, South Korea, Australia, India, etc.

In 2012 Open R&D was launched as an initiative to attract more industrial project leaders and partners to PRACE, so far 53 industrial projects have been awarded including large companies as well as SMEs.

PRACE awarded projects from major EU and international projects including FET Flagships (the Human Brain Project (HBP) and Graphene), large-scale instruments (ITER, LHC, ALMA, etc.) and European structured communities (ENES), as well as projects receiving ERC and Marie Curie grants.

The SSC provided the following recommendation to the PRACE Council:

"*PRACE 1.0 has been very successful in establishing a European HPC community, which has strengthened European science and competitiveness. Therefore, SSC strongly recommends that PRACE 2.0 continues to provide access to world-class HPC facilities in Europe at least at the level of PRACE 1.0.* "

PRACE SYSTEMS

Definition

The definition of PRACE 2.0 Tier0 systems is based on a combination of minimum capability (performance of the system) and minimum capacity (timewise slicing of the system):

- regarding the minimum capability of a Tier0, a new system should offer a peak performance at least 20% higher than the average of the peak performance of the existing PRACE systems;

- regarding the minimum capacity of a Tier0, a new system should provide at least 20% of the annually available cycles to PRACE peer review.

A more detailed information about PRACE 2.0 systems is provided in Annex 1.

Specification and Procurement

One of the strengths of PRACE is the provision of systems with a range of machine architectures, and it will be valuable to retain this in PRACE in the future. While procurements and operations of PRACE systems will be carried out autonomously and according to the governing rules/laws of the individual hosting institution, it is proposed that an annex to the Agreement for PRACE 2.0 will be produced (as in PRACE 1.0) where the HMs could indicate the main characteristics of the HPC architecture and a roadmap of deployment of the systems they plan to procure. Also, in PRACE 1.0 some of the PRACE benchmarks were used for selecting the systems and this could be done again in PRACE 2.0. It is proposed that SSC could iterate with PRACE 4IP to validate/update the content of the current PRACE benchmark suite.

MODEL FOR PRACE 2.0

It must be the aim for PRACE 2.0 to have a model which enables the same quality of HPC services as in PRACE 1.0, avoiding a reduction in capacity or capability. In other words, it is the goal to have PRACE systems that continue to offer the capability and capacity required to support leading-edge research, using state of the art technology trends.
As a starting point, we estimated the required budget per year to maintain the quality of HPC service to meet this goal, to be a total opex budget of €30 Million. By `fixing` the required budget, this allows the opex costs to be distributed amongst all members. The model uses a formula that takes into account past usage (PRACE 1.0) and GDP of each individual country. The model assumes that every member who uses resources has to contribute to the operational costs of the tier-0 systems. Clearly, the model allows us very quickly to check the overall and the individual affordability of the proposed business model for a particular total opex budget.

Another factor that has been built in to the model is the different opex costs of two different architectures (BlueGene type and X86 type). The HM contributions are calculated in the model using the same formula as above. This is then `converted` into millions of core hours each HM will provide to PRACE2.0 depending on the type of system, plus the fraction of the system that has to be provided to PRACE.

A key point to note about this model: an increase in the number of HMs does NOT increase the costs to other members, as the total costs are fixed. If a member becomes an HM, then their contribution is met through the provision of core hours as determined by the model.

The model indicates that with a budget of around €30 M pa, the resources available in PRACE 2.0 would be comparable with those in PRACE 1.0. Further details of the model and the costs are given in the paper and spreadsheet (Annex 2).

PRACE 2.0 LEGAL ENTITY

Legal advice from Bird & Bird has indicated that the AISBL can continue independently of the termination/completion of the agreement for phase 1 of PRACE; there is no automatic termination

of the AISBL. It is proposed for pragmatic reasons that we continue with PRACE AISBL as the legal entity for PRACE 2.0. The current Statutes can therefore remain in operation during the development phase of the Agreement for PRACE 2.0.

GOVERNANCE

As the current Statutes for PRACE can remain in operation, the current governance structures can also remain in place. However, with the development of a new agreement for PRACE 2.0, this gives Council the opportunity to improve the governance structures. It is possible to identify some issues to take into account:

- Transparency, simplicity, equality, and agility, in particular for decision-making.
- Voting rights reflecting the new PRACE 2.0 model
- Ensuring that the scientists and industrial users have a clear role.
- Leadership: putting in place a full time and onsite `managing director` (or whatever title is appropriate and agreed) of PRACE AISBL who will have responsibility for leading the operations of PRACE AISBL.

TIMETABLE

Date	Group/Meeting	Actions required
4th February 2015	Council	• Initial discussion and feedback on this proposal • Approval of preferred business model for PRACE 2.0
Feb – May 2015	SWG	• Further work to develop the business model based on feedback received
	Member states	• Engagement with EC
	SSC	• Scientific position paper about PRACE 1.0 achievements and the need of PRACE 2.0 (before end of February)
	Members	• Initial sharing of preferred business model with ministries • Feedback to SWG
Early May 2015	Extra Council	• Agreement on proposed business model
May - August 2015	Members	• Further interaction with ministries: discussion / agreement / feedback • Expressions of interest in being a member of PRACE 2.0 (as HM or GP)
	SWG (plus Bird & Bird)	• Development of the final version of the proposal, including any modification of Statutes required
End of Aug/ early September	Council	• Agreement on final PRACE 2.0 proposal
September 2015	Members	• Signature of PRACE 2.0 Agreement

ISSUES REQUIRING FURTHER DISCUSSION

The issues set out below require further iteration. Some initial thoughts and proposals are set out below. The SWG will continue to work on these, and feedback and input from Council on these matters is essential for that work to progress.

EC involvement and Industrial interactions

As explained above the model assumes at this time no EC contribution to operating costs. A major effort needs to be made to engage with the EC on their strategy and funding plans for HPC. The model proposed offers the EC the chance to gain significant leverage on PRACE member contributions. Any EC contribution to opex would enable the expansion of systems and services to European researchers.

PRACE needs to clarify its role in the development and support of the EU supply industry. One possible model is that in addition to a PRACE portfolio of very large production quality systems, 1-2 large scale (well beyond PCP) technology development, non-production platforms are also supported under the PRACE banner. This may well open up the possibility of getting sponsorship from the HPC supplier industries in Europe for co-developing machines for special applications, with DARPA-like funding from the EC. It would also allow the desired position of being able to retain the use of an open global procurement process for the production systems, while using a PCP-like process for the development platforms.

Efforts should also be made to get sponsorship/funding from HPC user industries for helping with use of HPC for better products and applications. For these industrial collaborations we should also seek funding from the EU and the relevant ministries (e.g. of economic affairs).

Operational costs

Since affordability is the major challenge for PRACE 2.0, it is essential that opex costs are kept to a minimum. We give here some indicative suggestions:

- Energy efficiency of the systems, (power consumption metered)
- User support – keep size of staff efficient; additional support can be given by the PRACE partners involved inside PRACE Implementation Projects.
- Maintenance costs with a maximum of approx. 10% of investment p.a.

There is a requirement for transparency about what is included in the costs of systems, and a robust audit process overseen by Council, building on the PRACE 1.0 KPMG evaluation process and lessons learned.

Role of Tier 1 Systems in PRACE 2.0

Tier 1 systems are a crucial part of the HPC ecosystem in Europe. There are a few reasons for this.

There is a risk that governments will not see their contributions to PRACE 2.0 as an addition to their investments in HPC, but as part of their existing overall budget. This potential reduction in budget available could damage national HPC provision, and the overall HPC ecosystem in Europe. A clear position needs to be developed on how PRACE fits in to the wider ecosystem of HPC, and in particular, the role Tier 1 systems can play in PRACE 2.0. In particular, the science and innovation benefits to be gained from access to PRACE systems needs to articulated in a way that will resonate with ministries.

Secondly, Tier 1 systems are being used as an important step-up to Tier-0 systems. Proven Tier 1 performance and scalability are required for computing time proposals to be taken into account for allocations on Tier 0 systems. Additionally, application enabling from Tier 1 to Tier 0 systems has

been (and still is) an important topic of the PRACE implementation projects. Lack of Tier 1 systems in certain countries may restrict certain researchers in their access to large systems, preventing them from doing excellent science. The risk that users may attempt to use PRACE systems for unsuitable computational projects in the absence of sufficient resources in their home country will be managed through the existing PRACE peer review processes, and as a result, no access to both Tier 1 and Tier 0 systems is the result. These issues will need further discussion and development for inclusion in the PRACE 2.0 agreement.

CONCLUSIONS AND SUMMARY

The proposed model offers an affordable way forward, allowing PRACE to continue to offer HPC resources for leading-edge research. It is recognised that there are issues that need to be addressed and developed further:

- Governance structures and voting rights
- Transparent definition of opex costs
- Process for auditing, and incentivising cost reductions and efficiency
- Interactions with and funding from EC
- Interactions with and funding from industry

These issues need to be addressed in parallel with the development of the formal Agreement for PRACE 2.0. PRACE members are encouraged to begin the process of discussing the proposed model and the financial and other implications with their relevant ministries as soon as possible, so that the target for signing the agreement in September can be met.

Annex 1

Definition of PRACE 2.0 Tier0 systems

Principles

The Tier0 qualification in PRACE 2.0 is based on the science supported, on the specific requirements of the scientific community and on the assumption that one solely architecture does not fit all science fields. This diversity of the systems, together with a centralized peer-reviewed open access based on scientific excellence are the main values achieved and offered by PRACE to Europe.

The proposed definition has 2 goals : on one side that such Tier0 systems will give the capabilities and the capacities required by scientists, thus enabling computational science that cannot be performed on smaller systems, and on the other side to increase inclusiveness of the infrastructure (in terms of hosting members and partners) and to enable co funding of the operations of Tier0 systems (by PRACE partners and possibly other funders including EC) towards the establishment of a persistent pan-European research infrastructure.

Like it was the case in PRACE 1.0 with the Annex III, future PRACE 2.0 Hosting Members will be invited to communicate in advance to PRACE Council any information about the architecture and the main characteristics (compute, storage) of the Tier0 they want to provide. Altogether future PRACE 2.0 Hosting Members will ensure that a complementary set of Tier0 systems, each scientifically productive in one or more domains and on a forward-looking architectural path, will be made available to scientific and industrial communities.

Starting from the assumption that the procurement of PRACE 2.0 Tier0 systems is under the total responsibility of the country providing it, future PRACE 2.0 Hosting Members are open to welcome criteria and recommendations elaborated by the SSC on the basis of the scientific communities needs.

Definition

The definition of PRACE 2.0 Tier0 systems is based on a combination of minimum capability (performance of the system) and minimum capacity (timewise slicing of the system):

- regarding the minimum capability of a Tier0, a new system should offer a peak performance at least 20% higher than the average of the peak performance of the existing PRACE systems;
- regarding the minimum capacity of a Tier0, a new system should provide at least 20% of the annually available cycles to PRACE peer review.

Extrapolation of the capability criterion

The following table shows the evolution of the peak performance of the PRACE 1.0 Tier0 systems made progressively available by the 4HM (based on the available information), and evaluates the minimum capability criterion as expressed previously:

Tier0 systems	Year	ΣPeak Perf Tier0 (in PFlops)	20 % * Avg (in PFlops)
JUGENE	2010	1	**1,2**
JUGENE, CURIE, Hermit, Fermi	2011	4	**1,6**
JUQUEEN, CURIE, Hermit, Fermi, SuperMUC, Marenostrum III	2012	15	**3**
JUQUEEN, CURIE, Hermit, Fermi, SuperMUC, Marenostrum III	2013	15	**3**
JUQUEEN, CURIE, Hornet Fermi, SuperMUC, Marenostrum III	2014	18	**3,6**
JUQUEEN, CURIE, Hornet Fermi, SuperMUC2, Marenostrum III	2015	21	**4,2**
JUQUEEN, CURIE, Hornet Marconi SuperMUC2, Marenostrum III	2016	29	**5,8**

Formattato: Italiano (Italia)

An extrapolation over the time of this trend is provided in the following graph:

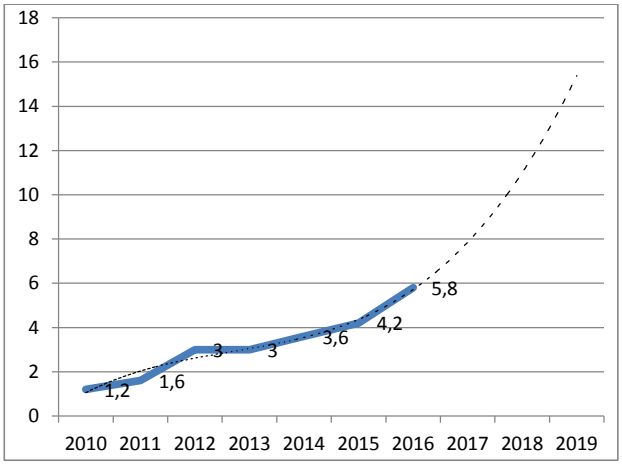

It shows that if we follow the trend of PRACE 1.0 this will lead to deploy in the future PRACE 2.0 Tier0 systems with a minimum peak performance of 4.2 PFlops in 2015, 5.8 PFlops in 2016 and around 8 PFlops in 2017, 11.5 PFlops in 2018 and 15 PFlops in 2019.

As this model could be distorted by the provision of very big Tier0 systems by some countries, the Council could re-evaluate this criterion on demand or on a yearly basis.

Annex 2

Analysis of PRACE 2.0 Business Model

The aim of the model is to maintain at minimum the same quality of HPC services including the expected evolution in capacity and capability based on state of the art technology trends as of PRACE 1.0. We try to estimate the future costs of such an infrastructure to analyse the affordability of the foreseen business model. The general principles for PRACE 2.0 stating that every member who uses resources has to contribute to the operational costs of the tier-0 systems. The OPEX costs are covered based on a formula that takes into account past usage (PRACE 1.0) and GDP of each individual country. Investment costs are covered by each individual HM solely.

We developed an excel sheet to be able to estimate the expected costs numerically. Some assumptions are made to achieve better understanding of the influence of the different parameters of the model. We tried to estimate the required budget of money per year to maintain the quality of HPC service as expected above. The needed individual contribution of each member is calculated by the distribution of the overall OPEX-costs by the individual percentage of the average of GDP and past usage. With a given budget we can check the overall and the individual affordability of the proposed business model very quickly.

The most important question for PRACE 2 with respect to science is: What will be the expected capability and capacity of the future PRACE tier-0 systems? In other words: What is the performance of the machines and how many core hours are available for PRACE projects? The dominating parameters of the calculation are the different OPEX-costs of the different architectures and the numbers of tier-0 systems available. We introduced two different systems in the excel sheet. System 1 might be a BlueGene type architecture with relatively low energy consumption and a high number of cores. System 2 is an X86 based system with only a few accelerators.

System 1 has the following characteristics:

- OPEX of core hour = 0.007 Euro
- Number of Cores = 300000
- Total number of Core Hours that can be provided by System 1 = 2.1 Billion annually. (assuming 80% utilization per year)

System 2 has the following characteristics:

- OPEX of core hour = 0.02 Euro
- Number of Cores = 100000
- Total number of Core Hours that can be provided by System 2 = 0.7 Billion annually. assuming 80% utilization per year.

Based on these values we can determine how many million core hours each HM will provide to PRACE2.0 if they had a system of Type 1 or a system of Type 2. The fraction of the system that has to be provided to PRACE for each HM is calculated in the excel sheet also.

Applying the model above with OPEX per year of 30 Million Euro, PRACE 2.0 can provide (as of 2015) 4.2 Billion core hours per year of System 1 and 1.5 Billion core hours per year of System 2. The numbers above are comparable to the resources provided in 2014 from PRACE 1.0 partners. For example CURIE in call 10 and 9 provided 220 to 230 Mcore hours. This is 450 Mcore hours per year. If you had 4 such systems this gives 1.8 Billion core hours i.e. close to 1.5 that the model provides for System of type 2.

In a similar way one can see that this is in average valid also for MareNostrum and Hornet. For a system of type 1 (Bluegene) we take Fermi as an example that provides around 400 per call i.e. 0.8 Billion per year. If PRACE had 4 such systems then the total amount of core hours available would be 3.2 Billion that is relatively close to 4.3 Billion that the model gives.

Summary:

The proposed business model implements the basic principles for PRACE 2 and provides enough flexibility to realise affordability, fairness and equal treatment of all members. With a given budget of 30 Mio per year for the overall OPEX costs PRACE 2 would be a useful continuation of the first phase.

		Member	Tier-0 Provider	GDP%	Usage%	50%/50% GDP/Usage	OPEX 50%/50% GDP/Usage	Contribution HM Mio Euro	Contribution GP Mio Euro	Cash-Back HM Mio Euro	Sum Mio Euro	System 1 core h (Mio)	System 2 core h (Mio)	System 1 fraction core h (Mio)	System 2 fraction core h (Mio)
Total Cost	60	Austria		2,14%	0,03%	1,09%	0,34		0,34						
OPEX	30	Belgium		2,63%	1,43%	2,03%	0,63		0,63						
		Bulgaria		0,28%	0,01%	0,15%	0,05		0,05						
		Cyprus		0,12%	0,31%	0,22%	0,07		0,07						
System 1	0,007	Czech Republic		1,06%	0,16%	0,61%	0,19		0,19						
Sys1 core	300000	Denmark		1,71%	1,68%	1,70%	0,53		0,53						
Sys1 core h	2102,4	Finland		1,35%	3,18%	2,27%	0,70		0,70						
		France	1	14,20%	18,79%	16,50%	5,12	5,12		3,34	8,46	1209	423	57,49%	60,36%
System 2	0,02	Germany	1	18,63%	20,81%	19,72%	6,13	6,13		3,99	10,11	1445	506	68,73%	72,16%
Sys2 core	100000	Greece		1,35%	0,20%	0,78%	0,24		0,24						
Sys2 core h	700,8	Hungary		0,68%	1,40%	1,04%	0,32		0,32						
		Ireland		1,14%	0,91%	1,03%	0,32		0,32						
		Israel		1,40%	0,08%	0,74%	0,23		0,23						
		Italy	1	10,94%	18,69%	14,82%	4,60	4,60		3,00	7,60	1086	380	51,63%	54,21%
		The Netherlands		4,19%	3,46%	3,83%	1,19		1,19						
		Norway		2,71%	1,07%	1,89%	0,59		0,59						
		Poland		2,66%	0,00%	1,33%	0,41		0,41						
		Portugal		1,15%	1,45%	1,30%	0,40		0,40						
		Serbia		0,21%	0,00%	0,11%	0,03		0,03						
		Slovenia		0,25%	0,00%	0,13%	0,04		0,04						
		Spain	1	7,19%	7,73%	7,46%	2,32	2,32		1,51	3,83	547	191	26,00%	27,30%
		Sweden		2,84%	1,76%	2,30%	0,71		0,71						
		Switzerland		3,43%	4,31%	3,87%	1,20		1,20						
		Turkey		4,28%	0,00%	2,14%	0,66		0,66						
		UK		13,45%	5,69%	9,57%	2,97		2,97						
				100%	93,15%	96,57%	30,00	18,17	11,83	11,83	30,00	4286	1500		
					Diff	3,43%		60,57%	39,43%						