SEVENTH FRAMEWORK PROGRAMME

Research Infrastructures

INFRA-2012-2.3.4 – Third Implementation Phase of the European High Performance Computing (HPC) service PRACE

PRACE-3IP

PRACE Third Implementation Phase Project

Grant Agreement Number: RI-312763

D2.1.4
Third Report on the joint PCP Pilot (Phase 3)

Final

Version: 1.0
Author(s): Fabio Affinito, CINECA
Paola Alberigo, CINECA
Date: 22.12.2017

Project and Deliverable Information Sheet

PRACE Project	**Project Ref. №: RI-312763**	
	Project Title: PRACE Third Implementation Phase Project	
	Project Web Site: http://www.prace-project.eu	
	Deliverable ID: D2.1.4	
	Deliverable Nature: Report	
	Deliverable Level: PU	**Contractual Date of Delivery:** 31/12/2017
		Actual Date of Delivery: 15/01/2018
	EC Project Officer: Leonardo Flores Añover	

* - The dissemination level are indicated as follows: **PU** – Public, **PP** – Restricted to other participants (including the Commission Services), **RE** – Restricted to a group specified by the consortium (including the Commission Services). **CO** – Confidential, only for members of the consortium (including the Commission Services).

Document Control Sheet

Document	**Title: Third Report on the joint PCP Pilot (Phase 3)**	
	ID: D2.1.4	
	Version: 1.0	**Status:** ***Final***
	Available at: http://www.prace-project.eu	
	Software Tool: Microsoft Word 2010	
	File(s): D2.1.4.docx	
Authorship	**Written by:**	Fabio Affinito, CINECA Paola Alberigo, CINECA
	Contributors:	Stephen Booth, EPCC Eric Boyer, GENCI Carlo Cavazzoni, CINECA Pekka Manninen, CSC Dirk Pleiter, FZJ Inigo Yenes, PRACE AISBL
	Reviewed by:	Mark Parsons, EPCC Veronica Teodor, FZJ
	Approved by:	MB/TB

Document Status Sheet

Version	Date	Status	Comments
0.1	04/06/2016	Draft	Outline of the final report document
0.2	01/12/2017		Added Stephen Booth contribution
0.3	13/12/2017		Added Carlo Cavazzoni and Dirk Pleiter contributions
0.4	15/12/2017		Added Inigo Yenes contribution
0.5	18/12/2017		Merged Eric Boyer contribution
0.7	21/12/2017		Added conclusions, removed the “Communication” section
1.0	22/12/2017		Final draft version for internal review

Document Keywords

Keywords:	PRACE, HPC, Research Infrastructure, Pre-Commercial Procurement, PCP, Methodology, Terms of References

Table of Contents

References and Applicable Documents

[1] http://www.prace-project.eu

[2] PRACE-3IP deliverable, D2.1.3 “Second Report on the joint PCP Pilot (Phase 2)”

[3] http://www.eocoe.eu/events/scientific-applications-towards-exascale

List of Acronyms and Abbreviations

AC	Assessment Committee
ADR	Alternative Dispute Resolution
ARM	Family of instruction set architectures for processors based on reduced instruction set computing (RISC) architecture developed by ARM Holdings
BCO	Benchmark Code Owners
BQCD	Berlin quantum chromodynamics program
CINECA	Consorzio Interuniversitario, the largest Italian computing centre (Italy)
CPU	Central Processing Unit
CSC	Finnish IT Centre for Science (Finland)
DLC	Direct Liquid Cooling
DoW	Description of Work
EC	European Commission
EPCC	Edinburgh Parallel Computing Centre (United Kingdom)
EU	European Union
FP7	7th Framework Program of the European Union for the funding of research and technological development in Europe
FPGA	Field-Programmable Gate Array
FZJ	Forschungszentrum Juelich (Germany)
GENCI	Grand Equipement National de Calcul Intensif (France)
GoP	Group of Procurers; a subset of PRACE-3IP partners committed to the PCP
GPGPU	General purpose computing on graphics processing units
HPC	High Performance Computing; Computing at a high performance level at any given time; often-used synonym with Supercomputing
HPDA	High Performance Data Analysis
HPL	High Performance Linpack benchmark

IDC	International Data Corporation, provider of market intelligence for the information technology
Intel KNL	Intel Xeon Phi “Knights Landing”: code name for the second generation MIC architecture product from Intel
Intel MIC	Intel Many Integrated Core Architecture or Intel MIC (branded as “Xeon Phi” since “Knights Corner” version)
IP	Intellectual Property
IPR	Intellectual Property Rights
JSC	Juelich Supercomputing Centre (FZJ, Germany)
LINPACK	A Benchmarks measure of a system's floating point computing power
LiU	Linkoping University (represented in PRACE by SNIC, Sweden)
MoU	Memorandum of Understanding
NDA	Non-Disclosure Agreement. Typically signed between vendors and customers working together on products prior to their general availability or announcement
NEMO	Nucleus for European Modelling of the Ocean
NPC	National H2020 Point of Contact
NRE	Non-Recurring Engineering
Pascal	Code name for the generation of NVIDIA GPGPUs released in 2016
PCP	Pre-Commercial Procurement
PE	Procuring Entity
PMO	Project Management Office
PPI	Public Procurement of Innovative solutions
PRACE	Partnership for Advanced Computing in Europe
PRACE aisbl	PRACE Association Internationale Sans But Lucrative
PRACE-3IP	PRACE 3rd Implementation Phase Project
RAPL	Running Average Power Limit
R&D	Research and Development
RI	Research Infrastructures
RUP	Responsabile Unico del Procedimento
SME	Small and Medium Enterprise
SoC	System on Chip
TCO	Total Cost of Ownership

TED	Tenders Electronic Daily (online version of the 'Supplement to the Official Journal of the European Union', dedicated to European public procurement)
Tier-0	Denotes the apex of a conceptual pyramid of HPC systems. In this context the Supercomputing Research Infrastructure would host the Tier-0 systems; national or topical HPC centres would constitute Tier-1
TRL	Technology Readiness Level
UEABS	Unified European Applications Benchmark Suite
USA	United States of America
WP	Work Package

Executive Summary

Deliverable Report D2.1.4 "Third Report on the joint PCP Pilot (Phase 3)" provides an update on the progress of the current PRACE Pre-Commercial Procurement (PCP) on "*Whole System Design for Energy Efficient HPC*". This is the last report in this series. The report explains how Phase II of the execution stage of the PCP was carried out and completed. Phase III is then described, including the timing and documentation used. Difficulties encountered during this final stage of the process are presented, with reference to previous European Commission (EC) recommendations. In particular, both the technical and procedural challenges that have been encountered during this part of the project are described. Finally, the lessons learned are presented, in order to take advantage from them for a future use of a PCP.

1 Introduction

In this deliverable, we describe the process of the Pre-Commercial Procurement during its Phase III, that started in June 2016 and ended on 31 December 2017 with the conclusion of the PRACE-3IP project. At the conclusion of Phase II, three contractors were admitted into Phase III of the PCP. According to Italian Procurement Law at least two contractors could have been admitted to the Phase III. At the end of the Phase II, the GoP decided that three contractors should be admitted into Phase III in place of two as previously planned. In order to allow the PCP with the three selected vendors to proceed a reduction of the budget was proposed and accepted by all the parties involved. As a consequence of this reduction, the design of the pilot systems was revised accordingly.

On 30 November 2017 the three admitted vendors (E4, Atos-Bull and Maxeler) submitted the End-Of-Phase III reports in which the results of the R&D services, all the required benchmarks and data about the installation and operations of the three pilot systems were documented. These reports were discussed and analysed by the Assessment Team appointed by the GoP. On 13 December 2017, the GoP had a final meeting to discuss whether the contractors performed according to the signed contracts and delivered the final decisions.

All the process described above proceeded smoothly without major issues.

2 Execution Phase III

As with Phase I and Phase II, Phase III started with the preparation of the tendering documents at the beginning of 2016. This process has already been presented in detail in deliverable D2.1.3 [2] available on the PRACE website [1].

The call for tender for Phase III opened on 14 March 2016 as defined by the PCP process plan in order to provide the tenderers with enough time to prepare the bids for the next phase and shorten at the same the gap between phases.

The Phase III call for tender closed on 5 May 2016 and the Procuring Entity received three bids from the following tenderers: BULL SAS, E4 Engineering, Maxeler Technologies.The Phase III selection process proceeded according to the timeline below:

- The call itself was issued on 14 March 2016 with a deadline of 05 May 2016
- All contractors awarded in PCP phase II also submitted a bid for Phase III.
- The evaluation of the proposals by the Assessment Committee took place on 25 May 2016
- The final assessment of the applications required clarification before the final ranking
- Requests for clarification were sent on 09 June 2016
- All answers arrived by 20 June 2016
- The Assessment Committee met again to analyse the answers on 27 June 2016
- The final ranking was agreed on 30 June 2016
- Communication of the public hearing for the opening of the Financial Offers took place on 29 June 2016
- The public hearing for the Opening of Financial Offers took place on 11 July 2016
- The GoP was notified on 14 July 2016
- The GoP met on 27 July 2016 to discuss the offers and the overall assessment results

Due to the close final ranking and differing but well justified views among the Assessment Committee team, the GoP decided to invite all three contractors into Phase III with a proposal for the reduction of configurations, but preserving the financial share devoted to R&D services

A proposal for the reduction of the final configurations was sent to the contractors, lead to the following budget:

- 1 M€ for Maxeler
- 2,47 M€ for E4
- 2,02 M€ for Bull-Atos

The contract was prepared and reviewed by the contractors who subsequently signed the contract and started the services for the PCP Phase III in the month of November 2016.

One important milestone of the execution of Phase III was the decision about the final hosting location of the three pilot systems. The GoP evaluated different possibilities and collected proposals among the PRACE partners, a final decision was take via a GoP vote on 14 February 2017.

As planned in the PCP execution plan, a set of PRACE experts performed visits to BULL (30 March – at R&D center in Clayes-sous-Bois, France), E4 (14 March – at E4 premises in Scandiano, Italy) and Maxeler (3 April – at Maxeler premises in London, UK) premises in March-April 2017 to monitor the progress of the services and potential issues. After the site visits, it was decided that each HPC site hosting a pilot system should take the responsibility to meet the contractor delivering the system to its site on a regular basis (at least once a month).

The BULL Pilot System FRIOUL was delivered on 14 April 2017 to CINES in Montpellier (France). Hardware installations including cabling, powering and pipe connections were finalized on 20 April 2017.

The software installation took longer than expected as well as the setup of the energy data collection infrastructure. Moreover, as for any system hosted in CINES, the PCP Pilot configuration had to follow strict security rules which was sometimes complex to setup with prototypes tools.

After, some early accesses, the PCP Pilot including first versions of BEO, HDEEViz and SLURM energy plugin became available for experimentation around mid-July 2017. This included PRACE-4IP WP7 extension activities.

The deployment of the Power Supply Unit with direct liquid cooling was performed in October 2017.

The openness of Bull-Atos pilot system was particularly successful, as PRACE activities were completed promptly. For example, the Center of Excellence EoCoE (energy oriented center of excellence) got access to the prototype and its tools, and a workshop was organised at CINES [3] .

E4 deployed the system (D.A.V.I.D.E.), an air cooled system, in their premises in April 2017, and validated it running the PCP benchmarks and HPL. They submitted HPL results to Top500 and Green500 for the June 2017 list. The system was then converted to water cooling during July and August 2017. Access to PRACE 4IP WP7 experts was granted for a small partition of the system to start porting applications in July 2017, when the system was still hosted at E4's premises. The system moved to CINECA's premises at the end of September 2017, and was made fully available to PRACE-4IP at the beginning of October 2017. Two main issues emerged during the execution. The first one was related to the late decision on the location of the system, E4 was forced to start the first validation in their premises, which was not planned at the beginning. This caused a delay of two month in the deployment of the system in the final location at CINECA. The second issue was related to the supplier of the direct liquid cooling components, CoolIT, who cancelled their support for the OpenPower architecture, in May 2017. E4 was then forced to find a different supplier. They subsequently found ASatek, a European company located in the Netherlands. This caused a second delay in the deployment of the system of about one month.

Maxeler deployed the final Phase III system in the JSC datacentre on the 27 October 2017. PRACE users gained access to the system immediately after its deployment.

3 The conclusion of Phase III

At the end of Phase III of the PCP the Contractors were obliged to submit an end-of-phase report. The structure of this report was defined in the Framework Agreement as follows:

- **Formal questions:** Replies to a set of formalised questions
- **Documentation of results:** Report which comprised a description of the technologies that have been developed, a documentation of the pilot system and a summary of the work that has been performed.
- **Installed and operational pilot system:** Report on the operational status of the pilot system including the required software stack.

- **Performance and energy efficiency results:** Report containing results on performance and execution times as well as the consumed energy for all provided PRACE applications.
- **Updated performance and energy consumption models:** Report comprising an updated version of the performance and energy consumption models for all four applications provided including an extrapolation to a 100 PFlop/s system.
- **Final market analysis:** a final analysis of the market.

All reports were received on time by 30 November 2017 and were forwarded the next day to the members of the final Assessment Committee. The Committee met in a remote meeting on 04 December 2017 with the task of discussing the reports, jointly assessing, whether the performance of the Contractors did meet the contractual obligations, and providing the GoP Committee with a recommendation. In its conclusions the Committee recalled the outcome of the monitoring visits, which were executed during Phase III and which allowed the PCP team to observe the progress of the R&D activities. The presented progress was assessed positively. The Committee furthermore noted that all Contractors did meet their obligations in delivering a pilot system and make this available for evaluation by PRACE experts. Furthermore, the Committee confirmed that the end-of-phase reports that were submitted did meet the contractual obligations. A small number of open questions were answered by the Contractors before a final meeting of the GoP Committee. The members of the Assessment Committee finally agreed that the contractors fulfilled their obligations and informed the GoP Committee accordingly. It did, however, also note the different levels of quality of the final reports. However, all were deemed compliant with the contractual obligations of the contractors.

The Assessment Committee recommended that the Lead Procurer perform an analysis of the financial reports with respect to the obligation that at least 50% of the budget was used by the contractors for R&D services. This analysis was required because from the financial sheets presented by the contractors the distinction between R&D and non-R&D work was not entirely clear (the contractors clearly found this confusing). As result of this analysis, it was concluded that all three contractors performed an amount of R&D work compliant with the terms of the PCP contract.

The GoP Committee met as planned on 13 December 2017. After discussing the report provided by the Assessment Committee, the GoP decided to ask for some clarifications via the Assessment Committee, related in particular to the final data received from the contractors and which was presented in the Assessment Committee report. The GoP met again on 20 December 2017 in Brussels to discuss the updated version of the report from the Assessment Committee. Also during this meeting some concerns about the detail of the PCP commitments from the vendors were discussed. The GoP discussed the proposal of delaying the final payments of the contractors, binding this operation to a complete satisfaction of the PCP objectives. The Lead Procurer and Project Manager were requested to clarify a short payment delay with the CINECA Administration and EC.

At time of the writing the final GoP decision on payment had not been made. However, as there is no concern that the three contractors have not delivered all of their commitment – but rather that some clarification on the document is required – the GoP have agreed to pay the three contractors should it not be possible within the confines for the FP7 Financial Guidelines to withhold a small proportion of the final payment.

4 Challenges and issues during Phase III

4.1 Technical

Challenges:

The PCP process has many phases and many formal decisions need to be taken. Delays in taking such decisions can clearly block or delay the execution of the PCP.

The constraints imposed by the national law of the procuring entity, required the physical presence in CINECA of the Assessment Committee members during the evaluation of the offers. This was inconvenient for the members and added little value to the process.

In general, the time for the assessment of the proposals and analysis of the deliverables was too short. This was particularly the case because the documents focussed not on existing products, but on new systems containing innovations requiring careful consideration.

Issues:

Allowing three vendors into the PCP Phase III increased the amount of activities for assessment, visits and monitoring by at least 50% with respect of what was planned. This was especially significant considering that the three vendors were in three different countries and the location of each final pilot system was in a different country as well.

It was not possible to transfer the property of the pilot systems during the execution of the R&D services, since this would have implied an acceptance of the system and an invoice, before the final assessment. This has caused some issues with the hosting procedures at the hosting sites.

A late decision in the final location of the systems caused a delay of two months in the delivery of the E4 pilot system, since they were forced to start the first validation at their premises. As a consequence, for the E4 pilot system, there was a limitation on the availability of the system for the PRACE-4IP project. However, on the positive side, this gave E4 the opportunity to validate two different cooling solutions and take measures useful for energy efficiency evaluation.

The Maxeler system was deployed at the very last stage of the PCP Phase III execution, and there was not enough time for the PRACE-4IP project, within the timeframe of the PCP, to test and evaluate the system. Anyhow the system will be made available to PRACE users after the end of PCP. PRACE experts will then have the opportunity to experiment with the pilot system.

4.2 Procedural

From the procedural point of view the following challenges and areas for improvement can be mentioned:

Hosting of prototypes:

To avoid delays in the planned schedule the hosting location of systems needs to be addressed at an early stage so it is settled before the last phase is launched. The Procurers need to agree on the sites where the prototypes will be installed and also under which conditions this will be done.

Duration of last execution phase

Taking into account that this is the last phase and since some extensions were already requested, the duration of this phase was crucial so the installation and testing could take place as planned. This is even more important when a whole system needs to be delivered as the different components need to be installed on time. In addition the procurers need to be able to access and test on site the prototypes.

Assessment of last phase

An important challenge faced by the Assessment Committee is the provision of an evaluation conciliating expectations committed to contractually by vendors for the last execution phase with reasonable expectations from the procurers for the actual outcome. Although the target innovation level may not fully be reached, the innovativeness of the results together with the R&D activity provided during that phase still needs to be assessed.

Furthermore, for further similar projects, we should take into account a stronger decision making process that will help to avoid delays on the whole procedure. A suggestion might be that the EC, perhaps via an expert representing their interests, become a member of the GoP (as a co-funder) in order to accelerate decision making.

Vendor engagement

There should be an obvious carrot for the winning vendor(s), e.g. a major contract as in the DARPA instrument of the Department of Energy in the United States.

5 Lessons learned

5.1 Technical

This project has explored the use of a Pre-Commercial-Procurement for the purposes of procuring/developing a highly energy efficient general-purpose HPC system. During the project we have learned lessons about both the use of the PCP instrument and about the technical challenges of energy efficient HPC systems.

The results of the PCP seem to demonstrate quite impressive energy efficiency for GPGPU accelerated computation in the current generation (when combined with appropriate cooling technology). This is not entirely general purpose HPC as significant software development effort is required to use GPGPUs effectively and this architecture is not appropriate for all applications. However, GPGPUs have now achieved significant market traction with a significant number of important codes supporting GPGPUs and the appropriate software skills are also widely available. This style of system should be seriously considered as an energy efficient option where the target workload is known in advance and can be evaluated for its suitability for GPGPUs. The Maxeler data-flow approach also shows promise. This approach also requires significant software development for each application but without the existing market traction enjoyed by GPGPUs may not be appropriate for all applications. If a machine is being procured for a limited set of applications the additional flexibility of FPGAs may be advantageous.

The PCP combines elements of a conventional hardware procurement with provision of funding for research and development. These two aspects are somewhat in conflict. Procurement is a commercial competition and needs to be run in a demonstrably fair and even-handed manner without any undue advantage being given to any of the potential vendors. Good practice in procurement is to have communication processes that ensure that equal information is provided to all competing parties. On the other hand, research is an essentially collaborative activity that benefits from free and ad-hoc communication. In this PCP though it did become clear that although this conflict existed, the impact was lessened because the aims of the PCP were closely aligned with one of the major issues facing the HPC market sector. As such the vendors already had a good understanding of the problem so there was much less domain specific knowledge that needed to be communicated. However, if the PCP mechanism is used for a more specialised set of requirements then a high degree of co-design between the end users and the vendors would be beneficial and the procurement process would have to be constructed carefully to ensure equal treatment of all competing vendors.

One recurring issue in the PCP was changes in delivery dates and road-maps of down-stream component supplies. Most of the energy budget of an HPC system is consumed by processors, memory and interconnect. All of these components require huge investments to develop and are usually provided by down-stream manufactures rather than the system vendors who engaged directly with the PCP. All of the vendors had to change and adapt their plans during the process because important components were delayed or cancelled. There is no indication that this situation is at all unusual – revisions in delivery dates and road-map changes appear to be the norm in the micro-electronics industry. This is not only a problem for a PCP type process. Conventional HPC procurements are also plagued by these phenomena and the commercial cost of HPC systems almost certainly reflects the wasted R&D costs that result from these types of change. However, there is no indication that a PCP process mitigates these problems in any way. If anything, it increases the exposure of the procuring entity. While the vendor will be able to recoup lost R&D costs from any future product in their portfolio the procuring entity requires a successful outcome for the final products of the PCP itself. There is also no indication that a PCP of this type has any significant influence over the down-stream vendors.

In contrast development of cooling systems does seem to be particularly well suited to a PCP type mechanism. It is essentially a mechanical design exercise that is clearly within the remit of the system vendor. The Bull product line is designed for the HPC space and was already largely water cooled. Some of the R&D effort within the PCP was used to increase the level of liquid cooling within this product line by extending it to the power supplies which previously had been air cooled. E4 also developed new liquid cooling solutions by modifying third-party compute nodes to convert them from air-cooled to liquid cooled. Cooling systems are largely independent of micro-electronics, software and other technologies driven by the down-stream vendors. Cooling systems have to be designed specifically for each product so there will be an additional cost associated with making a liquid cooled version of each product. Customers who wish to ensure that liquid cooling solutions are available can encourage this by helping to fund some of these development costs using a PCP type mechanism. This will be particularly useful if customers have specific requirements as part of a site energy strategy, such as wanting higher than normal inlet or outlet temperatures that make it easier to make use of waste heat. A PCP focused on cooling systems could then be used to enable some degree of co-design between system design and the design of the datacentre cooling infrastructure. Very innovative datacentre cooling infrastructure would run the risk of reducing the number of eligible vendors for subsequent procurements. It might be possible to mitigate this risk by also adding a small PCP phase to these procurements to help additional vendors develop compatible cooling solutions.

The aim of the PCP was to develop a general purpose highly energy efficient HPC system. In order to reflect these requirements, we evaluated the solutions using a representative set of benchmarks. In addition, the vendors were expected to use these benchmarks essentially unchanged – only a limited amount of changes to the software were permitted. While this would have been unproblematic for a conventional procurement, the extended time-scale of the PCP meant that the underlying codes evolved significantly during the process and insisting on the original reference versions of the code would have been counter-productive. It also became clear that it was very challenging to address energy efficiency in a totally general manner and that better results could be obtained by concentrating on specific problems or classes of problems. This is in part because energy efficiency is one of the main challenges facing the IT industry as a whole and the HPC sector in particular. Approaches that benefit the general case will often be applied irrespective of the existence of a PCP. When considering particular problems with specific requirements it might be better to evaluate based on a specified problem rather than a specified code version to allow for the evolution of the application codes during the process.

The energy efficiency of data-storage systems was within the scope of the PCP. The aim of the PCP was to support general purpose HPC work-loads. The current requirement of most HPC applications is to have the data-storage systems presented as Posix compliant parallel file-systems. None of the vendors chose to address this aspect of the system design in the early phases of the PCP. The file systems in the final pilot systems do reflect aspects of energy efficient design. For example, the selection of Solid State Disks over conventional magnetic media and to host the file-system directly on the compute nodes rather than provision an additional set of file-system servers. However, these are choices from within the space of conventional file-system design rather than new innovations which will result in new products. The evidence from this PCP therefore suggests that there is scope to influence storage system energy efficiency by setting appropriate procurement rules, but this should also be possible within a conventional procurement process and does not necessarily require the additional R&D support of a PCP.

The software and programming environment available on a machine is one of the key factors in making it successful. One factor that resulted in the substitution of Power8/GPU for ARM/GPU in the E4 solution was the relative lack of software support for HPC and GPGPU acceleration on ARM. Similarly, one of the strengths of the Maxeler system (compared to other FPGA based approaches) is the software environment.

The PRACE PCP did not place any restrictions on how energy efficiency was to be addressed. The systems were to be evaluated on total-system energy efficiency so vendors were free to address any and multiple aspects of the total system design. However, exploring multiple avenues simultaneously does run the risk that the different innovations may not be compatible at the end of the day. This was seen to some extent in Phase III where software tools that controlled energy efficiency by adjusting processor clock frequency gave promising results on conventional processors but were less promising when combined with the new KNL hardware. This risk could be reduced by focusing a PCP on a more tightly scoped area of research.

5.2 Legal

Legal framework

Due to the length of the procurement procedure the national legal framework may change at any point during any of the phases. With regard to this challenge, a regular check of any update of the applicable law is helpful to reduce risks in this regard.

Operational Plan for the Hosting of the prototypes

It is advisable to agree upfront on the different aspects related to the hosting of the prototypes. This can be done through an Operational Plan covering among others these points:

- Bearing of costs by the Hosting Site
- Access agreement
- System integration
- Service level from the Hosting site
- Security
- Decommissioning
- System evaluation
- Resource allocation
- Publication of results
- Insurance

Access terms for prototypes

In case this aspect is not addressed at an early stage the Procurers need to agree on the terms according to which the prototypes are accessed for testing and subsequent use. This aspect covers not only access from Procurers but also from final users in case this possibility has been foreseen (e.g. through a competitive procedure).

Property of prototypes

It is important to clarify if possible upfront when the transfer of property over the prototypes takes place, not only from the vendor to the procurers but also if any transfer takes place among the procurers.

Assessment of final phase

To facilitate the work of the Assessment Committee of the last phase it can be helpful to provide a scoring scheme for the different elements to be analyzed so there can be a weighting of the different aspects and also a way to check if a minimum level of completion has been reached. This would be particularly relevant if partial payment is foreseen.

6 Conclusions

In this document, we reported on the third phase of the PRACE-3IP project PCP, discussing the execution of Phase III and its conclusion. All the steps leading to the release of the contractors' deliverables and the installation of the three Pilot Systems were described. The most important challenges, with regard to the technical and procedural criteria were presented. This was important to inform any changes that should be considered in case of a new PCP in this and other highly technical areas. In order to better address these challenges, the authors proposed a list of "lessons learned" that will be usable for implementing a new PCP.

This PRACE PCP was the first use of PCP in the field of High Performance Computing. In many aspects, it was an important success story. However, many things could be improved with respect to this first experience. In this deliverable, we have provided many suggestions that we feel would be useful to take note of in order to improve the use of PCP in future projects.

However, our use of PCP was always planned to be an exploration of its efficacy in this domain. In this regard, and in terms of the finally delivered systems, the process we have followed has been a considerable success.