



**SEVENTH FRAMEWORK PROGRAMME
Research Infrastructures**

**INFRA-2007-2.2.2.1 - Preparatory phase for 'Computer and Data
Treatment' research infrastructures in the 2006 ESFRI Roadmap**



PRACE

Partnership for Advanced Computing in Europe

Grant Agreement Number: RI-211528

**D7.6.3
Evaluation Criteria and Acceptance Tests**

Final

Version: 1.3
Author(s): R. J. Blake, STFC Daresbury Laboratory
Date: 18.12.2009

Project and Deliverable Information Sheet

PRACE Project	Project Ref. №: RI-211528	
	Project Title: Partnership for Advanced Computing in Europe	
	Project Web Site: http://www.prace-project.eu	
	Deliverable ID: D7.6.3	
	Deliverable Nature: Report	
	Deliverable Level: PU	Contractual Date of Delivery: 31.12.2009
		Actual Date of Delivery: 18.12.2009
EC Project Officer: Maria Ramalho-Natario		

* - 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: Evaluation Criteria and Acceptance Tests	
	ID: 7.6.3	
	Version: 1.3	Status: Final
	Available at: http://www.prace-project.eu	
	Software Tool: Microsoft Word 2003	
	File(s): D7.6.3.doc	
Authorship	Written by:	Richard Blake
	Contributors:	UK – R.J. Blake, J.Nicholson France – J.-P. Nominé, F. Robin Germany – M. Stephan, S. Wesner Netherlands – P. Michielse Poland – N. Meyer, M. Zawadzki Italy – G. Erbacci
	Reviewed by:	Georgios Goumas, GRNET Dietmar Erwin, FZJ
	Approved by:	Technical Board

Document Status Sheet

Version	Date	Status	Comments
1.0	31.10.2009	Draft	Draft based on discussions at teleconferences.
1.1	22.11.2009	Draft	Input from teleconference and Norbert Meyer, M Stephan, S Wesner
1.2	30.11.2009	Draft	Input from G Erbacci, WP7 teleconference and vendors on pre-competitive procurement
1.3	18.12.2009	Final	Addresses comments from Reviewers

Document Keywords and Abstract

Keywords:	PRACE, HPC, Research Infrastructure
Abstract:	<p>Task 7.6 concerns the development of a Procurement Process Template to capture best practice in the purchase of large scale HPC systems. This document builds on Deliverable D7.6.1 which overviewed European procurement practices, reviewed a number of recent international procurements and commented on best practice and Deliverable D7.6.2 which developed a Pre-Qualification Questionnaire.</p> <p>Within this Deliverable we review more recent procurements by the Hosting and non-Hosting Partners and comment further on our experiences with the Negotiated Procedure and report initial views from a number of hardware vendors on the new pre-commercial procurement procedure. We overview best practice in the evaluation of:</p> <ul style="list-style-type: none"> • Pre-Qualification Questionnaires; • responses to technical requirements presented in Deliverable D7.5.2 and the performance of benchmarks discussed in Deliverable D6.2.2; • the assessment of risks discussed in Deliverable D7.4.2; and, • the overall evaluation of responses from vendors covering financial, corporate, technical and non-technical factors. <p>We discuss how the evaluation and assessment may vary between the assessment of novel architecture and general purpose systems. We conclude this Deliverable with comments on acceptance tests consistent with the specification of requirements.</p>

Copyright notices

© 2009 PRACE Consortium Partners. All rights reserved. This document is a project document of the PRACE project. All contents are reserved by default and may not be disclosed to third parties without the written consent of the PRACE partners, except as mandated by the European Commission contract RI-211528 for reviewing and dissemination purposes.

All trademarks and other rights on third party products mentioned in this document are acknowledged as own by the respective holders.

Table of Contents

Project and Deliverable Information Sheet	i
Document Control Sheet.....	i
Document Status Sheet	i
Document Keywords and Abstract.....	ii
Table of Contents	iii
List of Tables.....	iii
References and Applicable Documents	iv
References and Applicable Documents	iv
List of Acronyms and Abbreviations.....	iv
Executive Summary	1
1 Introduction	2
2 Update on Recent Procurements and Procurement Procedures	3
2.1 Procurement within Poland – N. Meyer	3
2.2 Procurement at Cineca – G. Erbacci	4
2.3 Procurement at HLRS – S. Wesner	5
2.4 Negotiated Procedure.....	5
2.5 Pre-commercial Procurement	6
3 Evaluation of Pre-Qualification Responses	8
4 Evaluation of Responses to the Statement of Requirements.....	9
4.1 Corporate Capabilities.....	9
4.2 Technical Capabilities.....	9
4.3 Benchmark Performance.....	11
4.4 Total Cost of Ownership	12
4.5 Risk Transfer	13
4.6 Added Value.....	14
4.7 Overall Evaluation.....	15
5 Acceptance Tests.....	16
6 Conclusions	19
Appendix A: Procurement in Poland by PSNC.....	20
Appendix B: Procurement by CINECA, Italy	21
Appendix C: Benchmark programmes	24

List of Tables

Table 1: Sample weightings for the assessment of technical capabilities of procured HPC systems ...	11
Table 2 : Sample weightings for the overall evaluation of procured HPC systems	15

References and Applicable Documents

- [1] <http://www.prace-project.eu>
- [2] PRACE – Grant Agreement Number: RI-211528 - Annex I: Description of Work
- [3] PRACE Deliverable D6.2.2 – Final Report on Applications Requirements
- [4] PRACE Deliverable D7.4.2 – Final Risk Register
- [5] PRACE Deliverable D7.5.2 – Technical Requirements for the second Petaflop/s system(s) in 2009/2010
- [6] PRACE Deliverable D7.6.1 – Procurement Strategy
- [7] PRACE Deliverable D7.6.2 – Pre-Qualification Questionnaire
- [8] PRACE Deliverable D7.1.2 – Final assessment of Petaflop/s systems to be installed in 2009/2010
- [9] PRACE Deliverable D7.1.3 – Final Assessment of Prototypes
- [10] PRACE Deliverable D6.3.1 – Report on available performance analysis and benchmark tools, representative benchmark
- [11] PRACE Deliverable D6.3.2 – Final Benchmark suite

List of Acronyms and Abbreviations

B	or Byte = 8 bits
CCRT	Centre de Calcul Recherche et Technologie du CEA (France)
CEA	Commissariat à l’Energie Atomique (represented in PRACE by GENCI, France)
CINECA	Consorzio Interuniversitario per il Calcolo Automatico dell’Italia Nord Orientale (Italy)
CoV	Coefficient of Variation
CPU	Central Processing Unit
DARPA	Defense Advanced Research Projects Agency (US)
DDR	Double Data Rate
ECC	Error-Correcting Code
ECMW	
EEA	European Economic Area
EPSRC	Engineering and Physical Sciences Research Council
ESFRI	European Strategy Forum on Research Infrastructures; created roadmap for pan-European Research Infrastructure
EU	European Union
ESP	Effective System Performance
Flop	Floating point operation (usually in 64-bit, i.e. DP)
Flop/s	Floating point operations (usually in 64-bit, i.e. DP) per second
GB	Giga (= 2^{30} ~ 10^9) Bytes (= 8 bits), also GByte
GByte/s	Giga (= 10^9) Bytes (= 8 bits) per second, also GB/s
GCS	GAUSS Centre for Supercomputing (Germany)
GHz	Giga (= 10^9) Hertz, frequency or clock cycles per second
HLRS	Höchstleistungsrechenzentrum, High Performance Computing Centre Stuttgart (Germany)

HPC	High Performance Computing; Computing at a high performance level at any given time; often used synonym with Supercomputing
I/O	Input/Output
ISC	International Supercomputing Conference, e.g. ISC'09 Hambourg June 24-25, 2009
ISO	International Standard Organisation
IT	Information Technology
ITT	Invitation to Tender
MPI	Message Passing Interface
NCF	Netherlands Computing Facilities Foundation (Netherlands)
OJEU	Official Journal (OJ) of the European Union (EU)
OS	Operating System
PAN	Polish Academy of Sciences
PB	Peta ($= 2^{50} \sim 10^{15}$) Bytes (= 8 bits), also PByte
PQQ	Pre-Qualification Questionnaire
PRACE	Partnership for Advanced Computing in Europe; Project Acronym
PSNC	Poznam Supercomputing and Networking Centre
RAM	Random Access Memory
RAS	Reliability, Availability and Serviceability
SMP	Symmetric Multi-Processing
SSP	Sustained System Performance
TCO	Total Cost of Ownership
TFlop/s	Tera ($= 10^{12}$) Floating point operations (usually in 64-bit, i.e. DP) per second, also TF/s
UPS	Uninterruptible Power Supply
WP	PRACE Work Package
WP4	PRACE Work Package 4 – Distributed system management
WP7	PRACE Work Package 7 – Petaflop/s Systems for 2009/2010
WTO	World Trade Organisation

Executive Summary

PRACE Task 7.6 has produced a Procurement Process Template to be used by the European Supercomputing Infrastructure, including the definition of a procurement strategy; the detailed implementation of which is addressed by other tasks within the work package.

PRACE Deliverable D7.6.1 [6] presented a brief overview of current European procurement procedures, recent procurements by the Principal Partners and discussed lessons learned in these and other non-European procurements. PRACE Deliverable D7.6.2 [7] developed a Pre-Qualification Questionnaire (PQQ) which supports the generation of a shortlist of potential suppliers for formal Invitation to Tender against the Statement of Requirements.

Within this Deliverable we review more recent procurements by the Hosting and non-Hosting Partners and comment further on procurement procedures. We overview best practice in:

- the evaluation of responses to the Pre-Qualification Questionnaire presented in Deliverable D7.6.2 [7];
- the evaluation of responses to technical requirements presented in Deliverable D7.5.2 [5] and the performance of benchmarks discussed in Deliverable D6.2.2 [3];
- the evaluation of risks discussed in Deliverable D7.4.2 [4]; and,
- the overall evaluation of responses from vendors covering financial, corporate, technical and non-technical factors.

We discuss how the evaluation and assessment varies between the assessment of novel architecture and general purpose systems. We conclude this Deliverable with a discussion of acceptance tests that are consistent with the specification of technical requirements.

1 Introduction

PRACE Task 7.6 has developed a Procurement Process Template to be used by the European Supercomputing Infrastructure, including the definition of a procurement strategy; the detailed implementation of which is addressed by other tasks within the Work Package. The elements of the Procurement Process Template include:

- Selection of an appropriate procurement process which complies with national and European regulations as discussed in Deliverable D7.6.1 [6].
- Shortlisting of credible suppliers. Deliverable D7.6.2 [7] developed a Pre-Qualification Questionnaire (PQQ) which supports the generation of a shortlist of potential suppliers.
- Evaluation of the responses to the PQQ and evaluation of the responses from the shortlisted suppliers to the full Statement of Requirements.
- Definition of Acceptance Tests consistent with the Requirements.

PRACE Deliverable D7.6.1 [6] presented a brief overview of European procurement procedures, procurements by the Principal Partners and discussed lessons learned in these and other non-European procurements. The key principles outlined in [6] can be applied both to the likely national procurements for the first Tier-0 systems as well as for a single European procurement by a future Research Infrastructure or equivalent. Feedback for the review of the PRACE project in March 2009 included:

Recommendation 24:

*'The procurement process and strategy has to be **further complemented with best practice examples of procurements in Europe and the negotiated procedures further elaborated.***

The suitability of the prototypes hardware and software to the specific user needs should be determined and used to inform the procurement process for the full scale systems.

The procurement should put equal emphasis on computer power as well as I/O, storage, visualisation, etc. The exact specification should be deduced from the benchmarking process on the pilot systems, of best suited user applications, i.e. the user applications that are best suited for top-tier HPC facilities.'

Within chapter 2 of this Deliverable we review more recent procurements by the Hosting and non-Hosting Partners and comment further on our experiences with the Negotiated Procedure and the potential role of the new Pre-Commercial Procurement procedure.

Deliverable D7.6.2 [7] developed a Pre-Qualification Questionnaire (PQQ) which supports the generation of a shortlist of potential suppliers for formal Invitation to Tender against the Statement of Requirements compliant with EU and national tender rules. Within chapters 3 to 5 of this Deliverable we overview best practice in:

- the evaluation of responses to Pre-Qualification Questionnaires produced in Deliverable D7.6.2 [7];
- the evaluation of responses to technical requirements presented in Deliverable D7.5.2 [5] and the performance of benchmarks discussed in Deliverable D6.2.2 [3];
- the evaluation of risks discussed in Deliverable D7.4.2 [4]; and,
- the overall evaluation of responses from vendors covering financial, corporate, technical and non-technical factors.

We discuss how the evaluation varies between the assessment of novel architecture and general purpose systems. We conclude this Deliverable with a discussion of acceptance tests that are consistent with the specification of technical requirements.

2 Update on Recent Procurements and Procurement Procedures

In Deliverable D7.6.1 [6] we reviewed a number of different nations' procurements of HPC systems and reported on best practice. These included:

- Procurement by NCF, The Netherlands
- Procurements within France
- Procurement by Jülich, Germany
- Procurement by Munich, Germany
- Procurement by Barcelona Supercomputing Centre, Spain
- Procurement by EPSRC, UK
- Procurement by ECMWF, UK
- NERSC – <http://www.nersc.gov/projects/procurements/NERSC6>

In Deliverable D7.6.1 [6] we reported on the CEA TERA 100 procurement and contract. TERA is the classified machine for defence applications, which corresponds to one branch of CEA's Supercomputing Complex, the others being CCRT – for research and industrial applications – and the forthcoming PRACE machine. These procurements will result in three machines / centres spread over two facilities: TERA (the generic facility for defence) and TGCC (the new facility designed to host next CCRT and future PRACE machines). By way of update, the final TERA 100 order is in process. The final purchase of a system was an option within the global contract which started off with a significantly sized R&D contract. The R&D phase was very successful, the outcome of this effort is related to the 'bullx' line of products announced in June 2010 at ISC (Köln University was the first commercial customer). The PRACE CEA WP7 prototype was also a precursor of this commercial series.

Since Deliverable D7.6.1 [6] as produced there have been further procurements embarked upon at PSNC, CINECA and at HLRS. The following sections summarise the processes followed.

2.1 Procurement within Poland – N. Meyer

About PSNC:

Poznan Supercomputing and Networking Center (PSNC) is affiliated with the Institute of Bioorganic Chemistry at the Polish Academy of Sciences (PAN). PSNC operates the scientific Internet in Poland and is amongst the top national HPC centres, being also a large R&D institution dealing with networking, grids, services & applications, security, etc.

Short description of the procurement procedures in Poland:

Any purchase above 14k EUR requires public institutions to follow the Public Procurement law which in essence requires that the desired purchase must be described in terms of its functionality, avoiding the naming of vendors or product names in a way that enables multiple bidders to compete for the contract. The specification then goes open to the public and its validity can be questioned by any interested party. After all specification issues are resolved the specification is closed, offers are collected, validated and scored according to the detailed rules described in the specification. The best offer is then chosen and again this can be questioned by interested parties. Finally, after resolving these issues the best offer is selected, contracts are signed and the purchase completed. For further details see Appendix A.

Selected examples of purchase criterions that led to successful purchases:

- Price: the least expensive offer wins (as long as it complies with the specification).
- Benchmarks: allow the procurer to specify the desired performance without naming a product. Examples include: SPEC.org tests (SPEC CPU2006), application-specific

benchmarks (e.g. the required number of Gaussian'03 jobs to be completed in a given amount of time), IO tests for storage systems, etc.

- Space utilisation: useful when small footprint solutions are desired due to space constraints (for example maximum. number of servers in a standard rack).
- Power efficiency: similar to above, more economic solutions are scored higher.
- Extended warranty: vendors providing support beyond the required number of years receive extra points.

Other interesting related notes:

- Sometimes it is useful to define the purchase as a gradual delivery over a number of years: it greatly simplifies the purchase procedures after the tender is completed and allows for upgrades when funds are released over a period of time (example: adding more computing nodes to a cluster every year).
- It is also possible to make a purchase from a selected vendor without following the Public Procurement law, however it must be very well argued (for example, only one product exists that complies with the requirements and the requirements cannot be changed in any way).
- Criteria can be assigned weights and used together.
- A good practice for larger purchases is to do extensive market research before publishing the tender (that includes meetings with vendors, testing products, analyzing prices, etc.).

Selected use cases:

- Gradual delivery of PC cluster elements in blade technology [tender published in July 2009, status for November 2009: offers collected and being scored]. Since many good blade products exist on the market, we decided to specify minimal requirements that enable most of the vendors (like minimum CPU performance, RAM size, maximum system footprint, etc.) and then give points for price (80%) and space utilization (20%).
- System upgrade of SGI Altix [October 2006]. Since an Altix SMP system can only be upgraded with original SGI parts that, for Poland, are distributed by a single SGI certified partner, PSNC was able to purchase CPU and memory modules without a tender procedure. However, an official statement from SGI and internal PSNC paperwork were necessary.
- Dedicated application system purchase [Oct 2005]. Here PSNC needed a system for Gaussian users. The specification included minimal requirements and a detailed benchmark procedure (including Gaussian input files and scripts) for each bidder to follow. Finally, we chose an AMD Opteron cluster that offered the highest price-performance ratio even though other solutions offered more computing power per CPU.
- service-oriented systems that are usually less complex are often purchased in a tender with a list of minimal requirements and a single scored criterion - the price.

2.2 Procurement at Cineca – G. Erbacci

In 2008 CINECA embarked on a procurement of a national HPC capability system to support the scientific and public research community. CINECA adopted a Competitive Dialogue Procedure and the procurement was structured in two phases addressing performance levels of 100 TFlops in 2009 and a petascale system in 2011. The process took about 18 months to conclude and has resulted in the purchase of an IBM p575 system in the first phase and an IBM Blue Gene Q system for the second phase.

The Competitive Dialogue procedure was selected because of the high technical specification of the hardware and the associated risks. The main technical requirements, of which in total there were more than one hundred, included peak performance, processing elements per node,

memory per node, MPI bandwidth and latency and storage system capacity. Information was gathered during SC'07 on technology roadmaps, software tools and applications and financial information. A PQQ and a Request-for-Proposals was formulated based on this input. The total budget was communicated to the vendors and they were invited to provide a financial quote for the phase 1 system and then specify a system for the second phase consistent with the remaining budget. The benchmark suite consisted of four applications

For further details please see Appendix B.

2.3 Procurement at HLRS – S. Wesner

HLRS runs a couple of high performance computing systems ranging from small systems mainly for technology evaluation (hardware and software), through medium sized systems driven by requirements to deliver cost-effective solutions also for commercial software vendor codes up to national computing facilities targeting the high-end user and delivering capability computing resources to them. The user community of HLRS has a focus on engineering in a broad sense but all kinds of research disciplines use the different levels of resources.

The procurement objective of the current system is the acquisition of the next generation of the national computing facility as part of the role of HLRS within the GAUSS Centre for Supercomputing (GCS). The delivery is anticipated in phases with a first system to be delivered in 2010 in the form of an intermediate system in order to allow in particular key users to familiarize themselves with a potential new computing environment and architecture. The installation dates of the major systems are broken down in two steps with an installation in 2011 and 2013. These dates had been coordinated with the other two GAUSS centres.

The initial phase of the procurement started with the publication of the tender at ted.europa.eu under the number 2009/S 144-211343. The procurement is following the Competitive Dialogue Procedure and accepted “request for participation” from any vendor until the 15.09.2009 outlining how the given criteria for participation (e.g. two reference installations of 100 TFlop/s systems) are met. All vendors that fulfilled the criteria for participation have been accepted and asked to respond to tender requirements. The tender requirements have been organised in three categories namely mandatory, important and desirable. The requirements cover different categories including performance and hardware topics, software requirements and collaboration aspects. Based on the presented offers and their analysis the competitive dialogue will be started in Q4/2009.

2.4 Negotiated Procedure

In Deliverable D7.6.1 [6] we reviewed the various EU procurement procedures and their suitability for acquiring different classes of systems ranging from research or novel architecture systems to general purpose systems. Within the negotiated procedure, a purchaser may select one or more potential bidders with whom to negotiate the terms of the contract. An advertisement in the OJEU is usually required but, in certain circumstances, described in the regulations, the contract does not have to be advertised in the OJEU. An example is when, for technical or artistic reasons or because of the protection of exclusive rights, the contract can only be carried out by a particular bidder.

Experience with the negotiated procedure amongst the PRACE partners is limited – the only system procured recently using this procedure is that at Jülich (JSC). The following requirements informed the selection of this procedure:

1. Technical limitations in the amount of available space and power / cooling capacity.

2. The maximisation of benefits for the investment in the existing 16 rack Blue Gene/P system - which was procured a year earlier - Jülich was explicitly looking for a solution that could be integrated seamlessly into the existing system infrastructure.

These requirements limited the number of suitable solutions to a very small number if not only one. In such a situation only a competitive dialogue or the negotiated procedure are reasonable procurement procedures. The advantage of the negotiated procedure is that a procurer can negotiate a separate contract with every bidder highly focused on the offered solution so the contract really meets the requirements. In this particular case, Jülich had a lot of experience with the Blue Gene/P system hence the contract with IBM did not need to include additional technical training. A contract for a different system by contrast would need to include training.

A further advantage of the negotiated procedure is that the finally selected solution(s) and contract(s) do not have to be the subject of a final tender; hence a procurer can choose the best fitting solution. In addition the whole procurement process – from the initial tender to the signing of the contract – can be completed in a shorter time period which can be critical, for example, when deadlines have to be considered.

2.5 Pre-commercial Procurement

In Deliverable D7.6.1 [6] we reviewed the new EU pre-commercial procurement which was introduced as a procedure in December 2007 with the intention of driving forward innovation in products and services to address major societal challenges.

http://ec.europa.eu/information_society/research/priv_invest/pcp/index_en.htm

The Commission organised a meeting in Brussels on 16th June 2009:

http://cordis.europa.eu/fp7/ict/pcp/events_en.html

The meeting reviewed examples of the use of the pre-commercial procurement procedure in a number of areas, mainly public sector service provision rather than high technology development projects. It is clear that the use of the procedure is very much in its infancy.

Informal discussions with technology vendors have raised the following issues:

- Non-European players have significant investments in Europe – difficult to exclude according to rules even if this was a good thing.
- European vendors need to compete in a global market.
- Lack of European only competition on the supply side of commodity components.
- Unwillingness of vendors to make their technology developments available in a competitive process.
- Europe should focus on the added value of hardware and software integration and delivering real performance in real applications to solve real problems.

In terms of specifying a procurement we would need to define clearly:

- the HPC requirements of these challenges that will not be met by the prototype/production/ novel architecture systems currently being evaluated;
- the innovations that are required to meet these requirements;
- the existence of sufficient procurement demand to support a substantial R&D activity that is likely to succeed; and,
- the European capability or capacity to potentially meet this need.

The general conclusion is that pre-commercial procurement is not well suited to the needs of developing unique systems such as supercomputers and that the research and development aspects are probably best met through a contract.

3 Evaluation of Pre-Qualification Responses

In Deliverable D7.6.2 [7] we developed a PQQ that requests Company Corporate Information in the following areas:

1. Company details and history.
2. Organisation and management.
3. Capabilities.
4. Financial.
5. Quality Management.
6. Supply Chain Management.
7. Legislative Compliance.

and information relating to the Specific Requirement including:

- Contact details.
- Staff qualifications and skills – resumes.
- Added value from other resources/ activities.
- Activities to be subcontracted.
- Financing – in particular of capital investment.
- Similar contracts undertaken elsewhere and evidence of performance.
- References from major international centres worldwide.

The weightings for the evaluation of a response to a PQQ will vary according to the type of system being procured, for example, whether access to reliable proven systems or the exploitation of novel architecture systems is the objective. In assessing the responses to the PQQ some of the criteria may be mandatory, for example if the supplier has evidence of legislative non-compliance, financial mal-practice or qualified accounts, less than minimum quality assurance or risk management certification, then they can be eliminated straight away.

Many of the evaluation criteria will need to be scored on a scale, possible scoring criteria might include:

- up to 3 points for economic and financial capacity reflecting relevant volume of sales;
- up to 2 points for technical accreditation standards above some minimum level;
- up to 2 points for maintenance/ support location and numbers; and,
- up to 3 points for reports from reference sites, e.g. 1 point for less than 3 reference sites, 2 points for 3-4 sites and 3 points for 5 or more sites.

For instance suppliers with say more than 5 points would then be invited into a formal procedure such as restricted or competitive dialogue or negotiated procedure.

4 Evaluation of Responses to the Statement of Requirements

As noted above the PQQ provides a mechanism for shortlisting vendors that are then invited to respond to a more detailed Statement-Of-Requirements, the responses to which are iterated through an Invitation-To-Negotiate prior to the selection of a preferred bidder with whom detailed negotiations are progressed to agree the contract.

If a PQQ has not been issued then the evaluation of the responses may well include an assessment that ensures compliance with mandatory requirements and a scoring of the ability to meet desirable requirements which is then incorporated through an appropriate weight into the overall score of the solution. If a PQQ has been issued then those vendors that have been shortlisted are assumed to have passed the Corporate Capability test and this category is not considered in the evaluation of the responses to the Statement of Requirements.

The evaluation of the Response to the Statement of Requirements needs to encompass the following aspects:

- Technical requirements.
- Benchmark performance.
- Total Cost of Ownership.
- Risk Transfer.
- Added Value.

In the following sub-sections we review best practice in each of these areas in terms of evaluation criteria and conclude this section with an overview of how to integrate these separable components into an overall evaluation framework.

4.1 Corporate Capabilities

Should these not have been considered through a PQQ process then the Corporate Capabilities should be scored as described in Section 3 and weighted appropriately as a category contributing to the overall score.

4.2 Technical Capabilities

Deliverable D7.5.2 [5] identified a broad range of **categories of Technical Requirements** which should be included as relevant to the particular procurement:

1. hardware including systems architecture and sizing;
2. I/O performance and global storage sizing, internal and external to the system;
3. post processing and visualisation;
4. software including operating system, management and programming environment;
5. operational requirements including installation constraints;
6. maintenance and support requirements;
7. training and documentation requirements; and,
8. delivery requirements.

The Deliverable avoided specifying final machine sizing by using minimum values and ratios, such as memory per compute node. This was designed to leave open the way future procurements are organised so, for example, it will allow a procurement to start with a fixed budget and seek to acquire the best performance for the available budget or seek the lowest price for a fixed performance. These requirements were provided on a per architecture basis.

Within the various categories of Technical Requirements there are different classes of requirements that focus on finer details of functionality and performance. These can be scored and weighted in the same way as the broad categories.

A subset of the technical requirements, relating to system sizing, includes specific requirement values, which unless otherwise stated, are minimum values to be met and allow the vendor to offer better values. Desirable elements give vendors the option of meeting them or not and to provide the opportunity for vendors to differentiate themselves from the competition.

Section 3.4 of Deliverable D7.5.2 [5] discussed the evaluation of vendor responses to the Technical Requirements. A quantitative method for comparing vendor submissions is to score the responses with a weighted points system. Responses can be scored as one of:

1. A fixed number of points if the requirement is met. The bid may be rejected if the requirement is not met.
2. A number of points per improved value over a base value up to a maximum.
3. A fixed number of points for the best performer and a reduced pro-rata value for worst.

The number of points are summed for each response and normalised so that the response with the highest number of points is assigned a value of 100. Normal rounding to a whole number is used.

For example:

Respondent	Technical Assessment Points scored	Technical Assessment Normalised points
A	34	81
B	42	100
C	28	67

This method allows other elements of the procurement, such as scoring benchmarking results, as well as non-technical elements such as capital costs and support costs to be combined into a final score as discussed at the end of Section 4.

The methodology outlined above can be applied to each of the individual Requirements and Desirable elements discussed in Deliverable D7.5.2 [5] contributing to a mark for each of the categories of technical requirement. In terms of the various categories we would propose the following weightings between the various desirable requirements with the maximum scores below summing to of order 100 points to enable the integration with other elements of the procurement.

In the following Table 1 we present sample weightings – we are not suggesting that these are mandatory in future PRACE procurements but it will be interesting to record future weightings in order to promote best practice. Some of the technical requirements may not be relevant to particular procurements, in this case then the category should be removed and the weighting of the other categories renormalized to give a total score of 100.

Technical Requirement	General Purpose System	Novel Architecture System
hardware including systems architecture and sizing;	50	75
I/O performance and global storage sizing, internal and external to the system'	5	2.5
post processing and visualisation;	5	2.5
software including operating system, management and programming environment;	10	5
operational requirements including installation constraints;	10	5
maintenance and support requirements;	20	10
training and documentation requirements; and,	mandatory	mandatory
delivery requirements	mandatory	mandatory
Total	100	100

Table 1: Sample weightings for the assessment of technical capabilities of procured HPC systems

The rationale behind the general purpose system is based to some extent on the weighting matching the proportion of funding spent on the requirement. Typically some 60% of the capital budget is spent on hardware for the compute, data and pre and post-processing systems, some 20% on maintenance over the period and some 20% on usability and manageability issues. Of course different technical specifications may well require different weightings allocated to different criteria.

The differences in weightings reflect the different nature of the systems with the general purpose systems requiring a balanced infrastructure supporting the whole lifecycle of applications development, running and post-processing in a highly available environment for a broad range of users. The novel architecture weightings reflect the emphasis on securing much higher levels of performance in a limited number of applications, and that the user community will have extensive experience requiring a less mature environment and are prepared to operate on systems with lower service levels.

4.3 Benchmark Performance

Deliverable D6.2.2 [3] discussed the applications and other software that will run on future European Petaflop/s systems. A recommended benchmark suite has been assembled by Task 2 in WP6 and documented in Deliverables D6.3.1 [10] and D6.3.2 [11] comprising synthetic benchmarks and representative application benchmarks. Within any specific procurement the user workload exploiting these and potentially new applications codes will need to be assessed and appropriate datasets constructed. This should give the vendors the opportunity to provide concrete performance and scalability data.

The performance of systems is usually assessed as a hierarchy of benchmarks (see Appendix C for references other than the PRACE applications benchmark suite), including:

- system component tests such as STREAM, P-SNAP (operating system noise test), SkAMPI, IOR, Metabench and Net Perf;

- kernels which run as serial through to full core count MPI on a node – tests memory bandwidth for various classes of algorithm;
- full applications such as those described in Deliverable D6.3.2 [11];
- composite tests measuring throughput such as:
 - SSP (sustained system performance) – which measures the geometric mean of processing rates for a number of applications multiplied by the number of cores in the system (for the highest core count runs);
 - ESP (effective system performance) – which measures the achieved job schedule against the best possible job schedule; and,
 - CoV (coefficient of variation) – which measures the variability of job run times.

The SSP provides a measure the mean Flop/s rate of applications integrated over time and thus takes account of hardware and software upgrades. The selected vendor is required to meet benchmark performance levels at acceptance and throughout the lifetime of the contract.

The overall mark for benchmark performance can be assessed using a variety of formulae capturing relative performance weighted by importance to the workload of the service. The final figure can again be renormalized linearly or non-linearly to a mark out of 100 reflecting the best performance. Deliverable D7.5.2 [5] suggests a pragmatic approach.

4.4 Total Cost of Ownership

Deliverables D7.1.3 [9] and D7.5.2 [5] presented an overview of the factors included in the Total Cost of Ownership (TCO) of HPC systems. The TCO for the system is an important figure that will need to be derived during a procurement process and matched to the available budget. Items that need to be included include:

- acquisition cost;
- maintenance cost;
- floor space requirement;
- power consumption for system plus cooling.

These will vary significantly depending on the nature of the system e.g. general purpose or novel architecture. There may well be site specific issues such as the need to integrate the system into the current mechanical and electrical infrastructure, systems/data infrastructure and the breadth and depth of skills required for systems management.

A typical service spends of order:

- 5-10 % of the budget on capital infrastructure (building, technical facilities, maintenance). It is quite difficult to capture the real cost of this in practice as the machine room may already exist, the capital cost may be depreciated only over the period of the current system's operations or over a building's natural lifetime, or the new system may need to fit into existing infrastructure. Given that most systems are designed to fit into 'normal' air-cooled or liquid cooled infrastructure these requirements can usually be categorised as mandatory – the system can fit into the proposed infrastructure (or the costs of adapting the infrastructure are marginal) or the costs of adapting the infrastructure are so large as to rule out the vendor.
- 10% of the budget on running costs such as electricity and systems management – again difficult to cost where the computer room may be sharing existing electrical and cooling infrastructure and the system management effort may be amortised over similar systems. The key metric here is the systems power and how that maps into usable flops.

- 65-70% of the budget on the system – supercomputer, related IT equipment and maintenance.
- 20% of the budget on scientific support.

The key issues that the vendor can respond to are:

- Cost of the equipment – includes system costs and infrastructure costs. The latter are either marginal in the sense that the system is being integrated into a current facility with specific wiring and cooling at marginal cost or requires a major upgrade to the facility which will make it uncompetitive if these are different to systems that fit.
- Cost of maintenance – depends on the reliability required of the system and the capabilities of the supplier. Usually included in the capital cost.
- System electricity consumption – ignores cooling requirements as this is usually defined by the infrastructure. The important metric here is the useable Flops/watt ratio as this really dictates the output from the system. There is no point having peak Flops for low energy consumption if the applications cannot take advantage of them.
- Ease of support and use - varies according to whether the system is a general user service on an established architecture (low) or on a novel architecture system (high).

The mandatory issues here are affordability covering both capital cost for the system and its recurrent budget. Best practice would point to specifying the overall budget in terms of budget for both capital and recurrent hence specifying the Total Cost of Ownership and optimising the most economically advantageous tender through appropriate weighting of the technical capability, TCO, performance, risk transfer and added value as discussed below.

4.5 Risk Transfer

Risks are assessed by their likelihood and impact and should clearly be managed by those parties best able to do so. The major areas of risks that involve the vendor were identified in Deliverable D7.4.2 [4] - Final Risk Register. These include:

- Risks that may prevent the system from becoming operational: a supplier ceasing to operate (2.2.1¹) or where a system fails to pass its acceptance tests (2.2.3).
- Risks that may delay the system operation: delays in the production process (2.3.1), delays in sub-contractors roadmaps (2.3.3).
- Risks that may limit the reliability, availability and serviceability of the system: lack of key functionality or performance e.g. global parallel file system (2.4.1) or power or cooling requirements may exceed expectations (2.4.3) or system may not be as reliable as needed (2.4.4).
- Risks associated with usage/ exploitation of the system: section errors in software and hardware (2.5.1) or section applications performing unexpectedly badly (2.5.2).

The various risks manifest themselves to the service provider in terms of potential demand risks (over demand/ under utilisation), price risks (need for extra infrastructure/ additional electricity costs/ inflation), timescale risks (failure to deliver the service on time to the community), performance risk (applications do not achieve the expected performance/ the system is less reliable than planned).

During the evaluation of tender responses the likelihood of these risks can be assessed for each vendor and marked and weighted in a manner similar to the evaluation of the technical requirements. The contract needs to incorporate various mitigation measures and appropriate penalties should the risks actually be realised. In terms of risk management, key risks are

¹ The numbers refer to sections in D7.4.2

typically held in a register and reviewed on an appropriate timescale. Deliverable D7.4.2 [4] proposed various actions to mitigate many of the risks which can involve the vendor, the service provider and the overall PRACE management at appropriate levels of escalation.

In summary, the relative importance of the various risks identified above will clearly vary according to the nature of the service. We suggest the following priorities in terms of addressing the risks:

1. Secure an operational system – availability and reliability of systems functions
2. Develop a functional system – all utilities, middleware and libraries available
3. Attain a reliable user service – environment consistent and reliable
4. Realize a productive system – performance for applications and overall workload

The contractual penalties should be negotiated in accordance with the relative importance of these priorities and focus on which party is best placed to manage the risk most cost-effectively.

4.6 Added Value

Added value assesses some of the broader tangible and intangible factors that the vendor can contribute to the overall service. These include direct added value to the service in terms of training, support for systems management and optimization in addition to facilitating interactions with vendor research labs and sites through to developing a joint business plan to further the aims of the centre. Clearly the range of possibilities is wide and procurers should select activities relevant to the service that they wish to support.

Developing the Service:

- Training and workshops on optimization.
- On-site specialists.
- Systems management user groups.
- Websites.

Developing scientific projects:

- Corporate R&D labs/ large-scale installation sites/ industrial partners/ public sector partners.
- Domain Scientific capabilities - data centric and compute.
- Technology Challenges – both systems and applications.
- Early access to technology – benchmarking service.
- Funding models – joint applications for Research grants, direct investment and hosting of staff.

These aspects would be rated more highly in the procurement of a novel architecture system.

Developing new products and services:

- In applications and systems software.
- In applications service provision.
- Working with current in-house service activities.
- Hosting R&D, support and marketing activities.
- Joint interactions with commercial software vendors.
- Attracting new/ other opportunities to locations.
- Using the site to host commercial systems.
- Outreach into higher education – addressing the skills agenda.
- International collaboration.

These aspects would be rated more highly in the procurement of a general purpose system.

Each of these categories should be converted into something of value to the customer: for example shared cost, shared profits, growing the community, enhanced scientific impact (e.g. quality and volume of publications), collaborations, or new projects. Scoring of these benefits will probably have to take place within the context of the broader strategy and business activities of the organisation and it is therefore difficult to be specific here. Added value is in the first instance a desirable activity and benefits should be assessed within that context.

4.7 Overall Evaluation

The overall evaluation inherits the markings in the various categories discussed above and then weights them to produce an overall valuation. In the following Table 2 we suggest weightings that may be appropriate to the categories for general purpose and novel architecture systems.

Area	General Purpose System	Novel Architecture System
Corporate Capabilities	10	5
Technical Capabilities	40	60
Benchmark performance	15	10
Total Cost of Ownership	15	10
Risk Transfer	10	5
Added Value	10	10
Total	100	100

Table 2 : Sample weightings for the overall evaluation of procured HPC systems

The category of Corporate Capabilities should be ignored if a PQQ has been issued. Likewise the Total Cost of Ownership category should be ignored if the capital and recurrent budgets have been specified in advance and the intent is to optimise the other categories. The remaining categories should have their weights renormalized as appropriate.

Care needs to be taken if the overall evaluation is combining many different levels of assessment as it may become possible to optimise a good score with unintended consequences. For this reason a sensitivity analysis should be run before the vendors are requested to respond.

5 Acceptance Tests

The acceptance tests need to address clearly the Technical Requirements summarised in Section 4 above and spelt out in section 5 of Deliverable D7.5.2 [5]. The different types of requirements can also be categorised as to whether they address the capacity of the system, its functionality or its performance.

In the following we briefly summarise the areas that should be subject to acceptance considerations and the test most typically used to assess the various requirements.

Capacity

The categories of most interest here are:

- maximum sustained Flops/s measured by Linpack;
- minimum applications memory per processing unit – tested through submitting jobs with a variable memory footprint;
- minimum global disk storage and maximum global file system partition sizes – measured by writing a single file of increasing size;
- minimum and maximum numbers of files available to an MPI job; and,
- archive and back-up file sizes measured by writing files of increasing size and number.

Functionality

Factors that should be tested here include:

- Standards – 32 bit and 64 bit arithmetic, ECC memory correction, UNIX-like POSIX operating systems functionality tested through a script of commands, support of multiple operating system images in different partitions and a security audit by an accredited body. Validated test scripts need to be developed here.
- Applications development environment for compilers, libraries, scripting languages, debugging and profiling tools. The environment needs to support a ‘module’ environment to meet the requirements of different applications. Most organisations have a collection of applications that will stress applications development environments. The vendors will need to demonstrate the tools working with selected users.
- On the systems side schedulers and monitoring diagnostics can be tested with throughput benchmarks and systems administration functions can be tested through shut-down and start-up exercises. Accounting and reporting utilities can be assessed whilst early users are accessing the system during the acceptance tests. Support for Grid applications can be assessed via tests of standard components discussed in Deliverables from Work Package 4 (WP4). Documentation is required for all of the major system utilities and can be viewed by inspection.

Performance

Should the performance of the system have been specified as part of the Statement-Of-Requirements then thought needs to be given within the contract as to how to deal with the situation that the measured performance does not meet the projected performance. This may result in the need to deliver more equipment which clearly has implications both in terms of the required infrastructure and running costs. There are many benchmarks of systems performance that can be included in the acceptance tests (see Appendix 3 for further details).

Factors to assess here include:

- Memory latency and bandwidth measured by STREAM.
- Point-to-point message passing latency, all-to-all bandwidth and barrier latency measured by SkaMPI.
- Peak read/ write bandwidths and latencies (MPI-IO, IOR, Metabench) for all processors for access to scratch and global file systems.
- Network connectivity and performance (GRIDftp).
- O/S memory/ CPU usage and jitter (measure with no applications), large page size efficiencies (test on selected benchmarks).
- Archive and backup performance (read/ write varying file sizes from global file system to archive).
- System resilience, start-up and shut-down tested during availability tests, support and maintenance arrangements can only really be tested in full production since the system is unlikely to fail during the acceptance period.
- Performance on appropriate applications from the PRACE Benchmark Suite (PRACE Deliverable D6.3.2 – Final Benchmark suite).

The analysis above should apply equally to systems with multiple components, for example, test and development systems, pre and post-processing systems and visualisation capabilities. There may be additional requirements to integrate the system into the current file system infrastructure and require interoperability not just between the components of the procured system but also with a range of other vendor offerings in particular clients for other file systems.

Clearly the specification of the technical requirements needs to be as specific as possible to meet the needs of the intended user community. Many of the requirements can be assessed by inspection or by using basic utilities or by running standard benchmark packages. Some of the requirements may need to migrate between explicit mandatory requirements and desirables depending on whether the system being procured is for a broad use community or for a specific application with a community prepared to soften its requirement for standards or up-front demonstrations of performance.

In addition to the baseline Technical Requirements which underpin the capabilities of the system service providers are particularly interested in the medium to long-term reliability, availability and serviceability of the system and users are interested not only in being able to routinely access the system but also in the performance of the system on their applications.

In terms of a timetable, the customer may wish a phased demonstration of the capabilities of the system. This may include:

- Factory test: comprising all hardware installation and assembly, burn in of all components, installation of software, implementation of production environment, low level tests (system power on and off), LINUX commands, monitoring, reset function, full configuration tests, benchmark performance.
- Systems delivery: delivered, installed, site specific integration prior to acceptance tests.
- Acceptance test periods – vary from 30 to 180 days. As noted above this may require a demonstration of the capacity, functionality and performance in addition to the various throughput tests discussed in section 4.3.

It should be noted that an aggressive approach to acceptance tests can be counter-productive resulting in very conservative proposals from the vendors which seek to minimise the risk that they take a long time to secure acceptance. A partnership approach may be more sensible for accepting novel architecture systems. It is always useful to specify reasonably flexible benchmark suites whereby marginal under-performance in one area may be balanced by

marginal over-performance in another area. Vendors should be given a reasonable period – say of order 180 days – to meet the most demanding availability tests which may run over a rolling 30 day period. The impact of delayed acceptance on the vendor’s bottom line should not be under-estimated – even for the larger vendors. The acceptance tests should also be consistent with the detail of the maintenance contract.

6 Conclusions

The aim of Task 7.6 was to develop a template for the procurement of Petascale systems. The task builds upon the various thematic studies focussing on the technical specification (Deliverable D7.5.2), infrastructure requirements and total cost of ownership (Deliverable D7.1.3) and risk register (Deliverable D7.4.2), incorporating their best practise with a review of lessons learned from recent EU and international procurements of HPC systems. Within the various Deliverables we have produced the various components of a Procurement Process Template and we have presented a synthesis of best-practice which provides the PRACE project with a sound basis for acquiring systems.

There will be an ongoing need to provide advice to partners on procurement best practise, learn lessons from new procurement exercises and refine the methodology to encompass the broader range of requirements and ever more complex technical solutions. Looking towards the future PRACE may well need to provide a portfolio of facilities ranging from general purpose systems that address the needs of a broad range of applications and users to novel architecture systems tailored to meet the needs of specific communities. This will require a programmatic approach to procurement addressing the short, medium and long-term requirements of the users.

Within PRACE Work Package 7 (WP7) we have focussed on the infrastructure and vendor aspects of the procurement of an HPC service. These systems need to be embedded in cost-effective service activities and some thought given to appropriate levels of service.

Appendix A: Procurement in Poland by PSNC

Norbert Meyer

Any purchase undertaken by an institution funded by the Polish government above a certain price (14 k Euro) has to be done in a strictly regulated way. This means in reality all major purchases, including bigger computing systems, have to follow the mentioned below rules:

- All internal orders for goods and services approved by department managers must afterwards be approved by the Public Procurement Officer.
- The Public Procurement Officer makes a decision whether a particular case is subject to tender procedures (*usually it is if the order is higher than 6 kEuro*).
- The procurement requirements have to be made publicly available on a dedicated web site and or delivered on demand to all interested parties (vendors).
- No specific vendor can be pointed out, except some special cases where the buyer can explicitly justify the reason, *e.g. the functionality requires only shared memory systems instead of distributed memory systems*.
- The specific solution or product cannot be required only, i.e. in addition the required functionality should be added (specific solutions can be mentioned as a reference so requirements stated like Xeon processor or similar are valid).
- The assessment criteria have to be clearly stated and the offer is assessed on scores computed on the basis of the criteria. There is a formal requirement that the price has to be included in the scoring, however the client can add other quantitative criteria that will be included in the score with different weights, e.g. warranty time or the max repair time.
- The offers from the potential vendors have to be made available to other competing vendors and they can look for any incompatibilities with the client requirements (except these parts marked as confidential).
- The assessment undertaken by the client/ buyer can be appealed against by the rejected vendor and a court decides whether the rejection was justified or not. This practice can make the procurement much longer, especially when the contract is a high priced one.
- The client can specify in the bid a period of time required for acceptance tests of the items delivered by the vendor prior to the payment. This in fact is not a regulation but rather a good practice.
- Within the tender procedures a commission evaluates the offers, announces the winner of the tender with whom the contract will be signed. If there are no objections within 7 days the contract can be signed.

Average time frames:

- 1) Tender announcement.
- 2) Answering additional questions, if any appear (+33 days).
- 3) Collecting offers (+7 days).
- 4) Evaluation process (about 7 days).
- 5) Announcing the winner.
- 6) Contract signature (+7 days, the 7 days are for any protests).
- 7) Delivery (the time depends on the time agreed in contract).

The whole time frame between the tender announcement and contract sign takes at least 8 weeks.

Appendix B: Procurement by CINECA, Italy

Giovanni Erbacci, CINECA, November 2009.

Introduction

CINECA provides computing resources and support to the scientific and public research community in Italy. This section describes the procurement of the national HPC capability system issued by CINECA. The procedure has taken about 18 months, starting with gathering input for a Request-for-Proposals (RfP) at SC07. The subsequent steps were the preparation of the RfP and the starting of an open European tender procedure in April 2008.

For the procurement CINECA adopted a *Competitive Dialogue Procedure*. A benchmark suite (low level and scientific applications) was provided and a set of specifications was issued to develop and find an optimal solution under the given budget.

The procurement was structured in two different Phases:

- Phase 1: Provision of a HPC system with a peak performance exceeding 100 TFlop/s to be delivered in 2009; and,
- Phase 2. Provision of a Petascale system in 2011.

The selected vendors produced the final offer in late 2008. In February 2009 the decision on the winning vendor was taken, and then the contract for the delivery was finalised and signed.

Information

The procurement was targeted for the selection and purchase of an HPC capability computing system and data storage equipment for the Italian scientific research community. The procurement was intended for the replacement of the previous IBM SP POWER5 (512 cores, 3.7 TFlop/s peak performance).

The overall timetable of the procurement process was:

- Information: November 2007 to December 2007.
- Preparing RfP: January 2008 to March 2008.
- Preparation of the benchmark suite: March 2008 to May 2008.
- European tender procedure: April 2008 to December 2008.
- Receipt of tenders: 1st step: July 2008; 2nd step: November 2008.
- Presentations by tenderers: 1st step: July 2008; 2nd step: October 2008.
- Judgement of proposals: 1st step July 2008, 2nd step December 2008.
- Benchmarking: May 2008 to November 2008.
- Reporting: January 2009.
- Communication to tenderers of decision: February 2009.
- Finalisation of contract with selected vendor: February 2009 to March 2009.

System selected for the first phase (2009):

- **Model:** IBM p-Series 575
- **Architecture:** IBM P6-575 Infiniband Cluster
- **Processor Type:** IBM POWER6, 4.7 GHz
- **Computing Cores:** 5376
- **Computing Nodes:** 168
- **RAM:** 21 TB (128 GB/node)
- **Internal Network:** Infiniband x4 DDR

- **Disk Space:** 1.2 PB
- **Peak Performance:** 101 TFlop/s.

System for the second phase (2011-12):

- **Model:** IBM BlueGene/Q
- **Peak Performance:** 1 PFlop/s.

The installation of the first phase System was completed in June/July 2009.

The Acceptance tests were completed in July and September 2009.

The Acceptance procedure included functional tests to verify on-site that the system satisfied the technical specifications and functional description (in terms of performance, memory, I/O, etc.). The benchmark suite was running through the acceptance tests in order to check and validate the values and the performances promised by the winning tenderer.

Procurement Procedure

CINECA choose a restricted procedure for the procurement, due to the high technical content of the hardware solution and the associated risks. The following process was put into action.

Gathering Information: During SC07, CINECA started to gather information on the state of the art HPC systems via *one-to-one* non-disclosure vendor meetings. All of the relevant vendors were informed, and appointments scheduled. A specific agenda of subjects to be addressed by the vendors during the meetings was prepared. The topics covered varied, from roadmaps to architecture and processor details, through software tools and applications, to financial information.

CINECA used the information gathered by the vendors to formulate a pre-qualification questionnaire and then to prepare the Request for Proposal (RfP). Requests for requirements gathered from the users were important to set up the final version of the RfP.

The pre-qualification step was based on the current financial standing and medium to long term financial viability of the vendors, associated with the capability to produce an adequate technical solution. Based on this information, four bidders passed the pre-qualification step and formed the short list and were invited to the negotiation, based on the *Competitive Dialogue Procedure*.

The RfP addresses two different phases within a fixed total budget.

- Phase 1: Provision of a HPC system with a peak performance exceeding 100 TFlop/s to be delivered in 2009; and,
- Phase 2: Provision of a Petascale system in 2011.

The total budget was communicated to the vendors. First, it was requested that the vendors make a financial quote for the system offered in Phase 1, then, with the remaining part of the total budget, it was requested that they offer a system for Phase 2, addressing a Petascale system. If the remaining part of the budget was not sufficient to provide an HPC system with a peak performance of 1 PFlop/s, it was requested that the vendors quote a system that cost the remainder of the budget.

The main requirements established in the RfP for the system in Phase 1 were:

- Peak performance of the whole HPC system exceeding 100 TFlop/s.
- Each compute node must be equipped with N PE, where $8 \leq N \leq 128$.
- RAM Memory: No less than 32 GB per node and at least 4 GB memory per core.
- MPI bandwidth between two compute nodes: at least 2 GByte/s.
- MPI bandwidth between two PE in a single compute node: at least 1 GByte/s.

- MPI latency between two compute nodes: less than 6 micro sec.
- MPI latency between two PE in a single compute node: less than 2 micro sec.
- Interconnection network with no latency or bandwidth degradations for applications requiring up to 512 PEs in “non blocking” mode.
- Storage subsystem for a total capacity of 1.2 PB.

Overall, there were more than 100 requirements (covering technical and performance aspects of the system, but also financial aspects of the company, future technologies and roadmaps), some of them mandatory, others evaluated as a function of a given weight, proportional to their importance. The weights have been fixed by CINECA and the vendors were aware of the given weights before submitting their proposals. Each vendor had to pass the mandatory requirements.

- A first proposal was presented by the four vendors in June 2008 and after that step, vendors were invited to run the benchmark suite.
- In July 2008, in depth meetings were conducted between the vendors and CINECA in order to refine their offers and to produce a Best and Final Offer for November 2008. For that date each vendor had to produce all of the benchmark results.
- The final offers were reviewed, scored and ranked against each other.

Benchmark suite

The benchmark suite was composed of four computational applications representative of the application workload of CINECA. The applications were tuned for scaling to a large number of processors, and relevant input sets were set up. To complete the suite, some synthetic benchmarks were added.

The benchmarks were set-up and run by the vendors in the period July – October 2008, assisted by CINECA staff. Each vendor had the opportunity to send to CINECA improved benchmark results up to November 2008.

Procurement commission

All of the procurement activity was addressed by a commission issued by the CINECA BoD and composed of six members, four of them external to CINECA, and appointed by the BoD, and two internal appointments: the Director of CINECA and the Director of the Systems and Technologies Department.

Based on the advice of the procurement commission, the CINECA BoD selected the vendor which won the tender and invited him to finalise the contract in February 2009.

Appendix C: Benchmark programmes

Benchmark	Functionality	Reference
STREAM	Sustainable Memory Bandwidth in High Performance Computers	http://www.cs.virginia.edu/stream/
P-SNAP	Operating system noise	http://www3.lanl.gov/pal/software/psnap/
SkaMPI	MPI implementations	http://liinwww.ira.uka.de/~skampi/
IOR	Parallel file systems	http://sourceforge.net/projects/ior-sio/
MPI-IO	Parallel file systems	http://public.lanl.gov/jnunez/benchmarks/mpiioetest.htm
Metabench	Synthetic and practical benchmarks	http://www.7byte.com/index.php?page=metabench
Linpack	The benchmark for HPC systems performance	http://www.netlib.org/benchmark/hpl/
NAS	Kernels and pseudo applications for CFD applications	http://www.nas.nasa.gov/Resources/Software/npb.html
HPCC	HPC Challenge benchmark – component and kernel tests	http://icl.cs.utk.edu/hpcc/
SSP – Sustained System Performance	A methodology for evaluating mean flop rate of applications integrated over time – workload performance metric	http://escholarship.org/uc/item/4f5621q9
Effective System Performance	Methodology for measuring system performance in a real world operational environment.	http://www.nersc.gov/projects/esp.php
Coefficient of Variation	Methodology to measure system variability	http://www.nersc.gov/projects/esp.php
Netperf	Network performance	http://www.uk.freebsd.org/projects/netperf/index.html
Grid ftp	Wide area file transfer between distributed systems	http://www.teragrid.org/userinfo/data/gridftp.php