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

D7.6.1
Procurement Strategy

Final

Version: 1.2
Author(s): R. J. Blake, STFC Daresbury Laboratory
Date: 17.12.2008

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: D7.6.1	
	Deliverable Nature: Report	
	Deliverable Level: PUBLIC	**Contractual Date of Delivery:** 31 / 12/ 2008
		Actual Date of Delivery: 31 / 12 / 2008
	EC Project Officer: Maria Ramalho-Natario	

Document Control Sheet

Document	**Title: Procurement Strategy**	
	ID: D7.6.1	
	Version: 0.2	**Status:** Draft
	Available at: http://www.prace-project.eu	
	Software Tool: Microsoft Word 2003	
	File(s): D7.6.1.doc	
Authorship	**Written by:**	R. J. Blake
	Contributors:	UK – R.J. Blake / J.Nicholson France – JP Nominé / F. Robin Germany – M. Stephan / S. Wesner Spain – E. Quingles/ S. Girona Netherlands – P. Michielse / A.B. Sa Dias
	Reviewed by:	Orlando Rivera, LRZ Dietmar Erwin(PMO), FZJ
	Approved by:	Technical Board

Document Status Sheet

Version	Date	Status	Comments
0.1	24/11/2008	Draft	For discussion
0.2	03/12/2008	Draft – For review at conference call 03/12/2008	Input from ECMWF and CEA.
1.0	05/12/2008	Draft version for internal review	Reformatting plus additional input for CEA, FZJ and BSC.
1.1	09/12/2008	Version for Internal Review	Input form LRZ Comments from WP7 participants
1.2	17/12/2008	Final Version.	Incorporating internal comments

Document Keywords and Abstract

Keywords:	PRACE, HPC, Research Infrastructure
Abstract:	This document builds on Deliverable 2.3.1 which summarised the procurement process on an abstract level in order to provide an overview of the different stages performed starting from a statement of requirements, development of a business plan, specification of a tender, the bidding process, development of the contract, acceptance and evaluation of the installed system. This deliverable overviews European procurement practices, summarises recent procurements by the Principal Partners within PRACE – reviews a number of recent international procurements and comments on best practice.

Table of Contents

List of Tables

References and Applicable Documents

[1] http://www.prace-project.eu
[2] PRACE Deliverable 2.3.1 – Discussion on Procurement Strategy
[3] PRACE Deliverable 2.1.1 – Report on Options for a Legal Entity
[4] PRACE Deliverable 2.2.1 - Overview of the Governance Structure of Proposed PRACE Entities
[5] PRACE Deliverable 2.7.3 – Report on Selection of First Production Systems
[6] Pre-competitive procurement: http://ec.europa.eu/information_society/research/priv_invest/pcp/index_en.htm
[7] NERSC Procurement Request for Proposals: http://www.nersc.gov/projects/procurements/NERSC6

List of Acronyms and Abbreviations

ESFRI	European Strategy Forum on Research Infrastructures; created roadmap for pan-European Research Infrastructure.
HET	High Performance Computing in Europe Taskforce. Taskforce by representatives from European HPC community to shape the European HPC Research Infrastructure. Produced the scientific case and valuable groundwork for the PRACE project.
HPC	High Performance Computing; Computing at a high performance level at any given time; often used synonym with Supercomputing.
HPC-Europa	Consortium of six leading (HPC) infrastructures and five centres of excellence providing transnational access; EU project.
ITER	Joint international research and development project that aims to demonstrate the scientific and technical feasibility of fusion power. Also used as the name for the reactor.
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.
PRACE	Partnership for Advanced Computing in Europe; Project Acronym.
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.
TCO	Total Cost of Ownership. Includes the costs (personnel, power, cooling, maintenance, ...) in addition to the purchase cost of a system.

Executive Summary

PRACE Task 7.6 intends to develop a Procurement Process Template to be used by the European Supercomputing Infrastructure, including the definition of a procurement strategy; the detailed implementation of which would be addressed though other tasks within the work package. Task 2.3 reviewed the procurement process as a generic process from an abstract viewpoint identifying stages and stakeholders. In the absence of progress in defining the operational model for PRACE (operational versus cycles) we present a brief overview of current European procurement procedures, recent procurements by the Principal Partners and discuss lessons learned in these and other overseas procurements which should inform future tasks within this work package. The principles outlined in this document 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 European Research Infrastructure or equivalent.

1 Introduction and background

1.1 Objectives

Task 7.6 is labelled in the work programme as **"Task 7.6: Procurement process template".**

The procurement process to be used by the European HPC Research Infrastructure will be defined, including:

- The definition of the procurement strategy, e.g. separate tenders for hardware, support and management, scientific support or single tender for the whole service.
- Pre-qualification questionnaire identifying organisations and/or suppliers with appropriate technical capabilities and financial standing and/or request for information (RFI).
- Technical specification as defined in Task 7.5.
- Vendor selection criteria such as quality derived from the responses to the technical requirements to be combined with benchmarking results, delivery, expertise and organisational capability, investment costs, total cost of ownership, price performance.
- Acceptance criteria and verification.

The following deliverables document the work:

D7.6.1 Procurement strategy (M12)
D7.6.2 Pre-qualification questionnaire (M18)
D7.6.3 Evaluation criteria and acceptance tests (M24)

This deliverable refers to work undertaken within Task 7.6 to deliver D7.6.1.

1.2 Dependencies and Related Deliverables

There are three dependencies and related deliverables that should guide the work for Deliverable 7.6.1.

- Deliverable 2.3.1 which defined a generic system procurement strategy,
- Deliverable 2.1.1 which looked at options for the legal form for PRACE,
- Deliverable 2.2.1 which looked at the governing structure.

1.2.1 Deliverable 2.3.1 - Document on Procurement Strategy

It was recognised in the summer of 2008 that there was a dependency between Task 7.6.1 and Task 2.3.1. The definition of Task 2.3.1: **Specification of funding and usage strategies** includes:

Principles for the procurement of production systems (HW, SW, Services) will be defined. Especially the pre-commercial procurement will be evaluated.

Deliverable 2.3.1 focussed on the principles of the procurement process based on the best practice of the involved institutions collected over the last years in the procurement of a variety of systems. It included an overview of stages in the overall procurement process and

summarised the activities envisaged within these various stages, the expected results and the roles and responsibilities of stakeholders. The deliverable aimed at a definition of the procurement strategy from an abstract viewpoint for hardware system, to be instantiated and detailed for the procurement of a specific system. It aimed at describing the guiding principles and it was not intended to be used as a step-by-step guide for the procurement of a Tier-0 system as it could not provide the necessary detail as at it had not yet been decided how responsibilities would be distributed between PRACE and its member organisations.

Deliverable 2.3.1 was reviewed by the Commission external reviewers with the following feedback presented at the WP7 teleconference in September 2007: **Recommendation 5: The procurement process and strategy has to be further complemented with best practice examples of HPC procurements in Europe, the USA and Japan and the negotiated procedures further elaborated. Those have to be documented in the most related follow up deliverable of WP2 (or WP7).**

1.2.2 Deliverable 2.1.1 - Report on options for a legal entity

The different options ranging from international treaties, European legal forms (including the newly proposed European Research Infrastructure) to national law were described and compared. The basis for this comparison were criteria, which were selected to meet the specific needs and requirements of the PRACE Research Infrastructure. References to 'procurement (strategy)' within the document included.

"3.9 Flexible for Usage": The legal form should be able to accommodate a possible transition of the PRACE organization from the cycles model to the operator model. This clearly affects who procures the systems and hence the legal constraints governing procurement policy.

Under the **Cycles Model** the hosting partners would be delegated the responsibility for designing, acquiring and operating the facilities and the supercomputers, taking into consideration the strategic needs of PRACE. The PRACE entity would perform processes for service definition, high level requirement analysis and contracting with the hosting partner.

Under the **Operator Model** the PRACE entity would be responsible for procurement and would own, house and operate the systems from its own budget.

The model used, or combination of models used will depend ultimately on a PRACE Principal Partners committee/management board decision. One possible situation is that the organisation would begin based on the cycles model and evolve in the general direction of the operator model in a period of 2-5 years.

The **Operator** and **Cycles Models** at this point conflate the issues of funding for the PRACE infrastructure with supply of the infrastructure. One could conceive of a situation where PRACE could pool its funding but still subcontract the procurement of systems to specific sites. There are good reasons for this because in a procurement one is seldom acquiring a system in isolation, the system will need to be housed in specific accommodation infrastructure in terms of building, power and cooling, and may need to be integrated with other storage, back-up and networking. Maintenance needs to be included in the procurement to ensure value for money. Supercomputer centres have a great deal of experience in leading procurement exercises (often subject to external parties specifying the requirements) and it may not be cost-effective to duplicate this level of expertise in a separate PRACE organisation and it may be sensible to avoid issues such as treatment of capital depreciation, insurance and software licensing.

"3.10 Tax Exemption": Tax exemption is important from a financial point of view. PRACE will employ a continuous procurement policy, and an annual budget of at least 170M Euros is foreseen. The legal entity needs to be able to establish tax efficient vehicles for these individual procurements. One strategic issue that may need to be resolved is the payment of VAT which may not be payable if commercial work on the system does not exceed 10% or if a significant proportion of the work is to support medical research.

"3.11 Flexible Procurement": It may be convenient for the PRACE entity, in certain cases, to engage in the **procurement of immature technologies** which are seen as key to the future of European HPC. This **pre-competitive procurement** (as mentioned in the project DoW) **is not possible** under present European and national procurement laws. **So it may be appropriate that PRACE be exempt from national and European procurement rules for certain situations.** This does not affect the fact that PRACE should always ensure that it follows appropriate rules in all procurements in order to **get the best value**.

The comment on "pre-competitive procurement" above is probably incorrect in general. The communication from the Commission in December 2007 (see discussion in section 2.6 and web link) was aimed at drawing the attention of the Member States to the existing but underutilised opportunity of pre-commercial procurement.

The Executive Summary and Analysis of Options in Section 5 concludes:

- The establishment of an international legal entity for the PRACE RI based on international treaties within the project duration of two years is unlikely because of the lengthy and complicated international negotiation process.
- European Research Infrastructure (ERI) by design seems quite suitable for research infrastructures like PRACE. All the special requirements of a distributed pan European research infrastructure are taken into account and a decision will be need ultimately based on the final valid Council regulation. However, the availability of the new ERI legal framework in time for PRACE is not yet assured, as the proposal still has to be approved by the European Council. The ERI is an international body within the meaning of Article 15, point (c), of Directive 2004/18/EC. Under the ERI, PRACE will not have to follow EU procurement law. This does not affect the fact that PRACE should always ensure that it follows appropriate rules in all procurements in order to get the best value.
- In the case that an ERI will be not available for the PRACE Research Infrastructure, a national legal from has to be selected. Different national legal forms from the Principal Partners countries as well as Belgium were considered.

Implications for D7.6.1:

- Ensure we assess pre-competitive procurement in the evaluation of procedures.
- Ensure we make the procurement strategy portable from cycles to operations model (for systems).
- The only viable Legal Form for early procurements is likely to be a national one.

1.2.3 Deliverable 2.2.1 – Overview of the Governance Structure of Proposed PRACE Entities

These must clearly match the legal entity and support a number of functions. Under the Cycles model PRACE will have to be able to:

- Delegate to the Partners or the Tier-0 hosting sites or organisations the responsibility to perform several of the tasks required.
- Implement and manage contracts held with Tier-0 hosting organisations.

Under the Operator model PRACE will also have to be able to:

- **Procure** Tier-0 systems
- **Procure the accommodation** for the Tier-0 systems
- **Operate** the Tier-0 systems
- **Provide support** to the Tier-0 System users .

Implications for D7.6.1:

- Ensure we assess pre-competitive procurement in the evaluation.
- Do not include procurement of site within this deliverable. This may be required if PRACE moves from the cycles model to an operator model.

1.3 Summary

So, overall, in terms of PRACE Deliverable D7.6.1 the following points need to be taken into account:

- Limit the procurement discussion to issues concerned with the procurement of individual systems. PRACE Deliverable D2.7.3 'Report on Selection of First Production Systems', which is being completed at the same time as this deliverable, states in essence that the initial procurements will most likely be national ones.
- Discuss the strengths and weaknesses of the various EU procurement procedures including pre-competitive procurement as the most likely vehicle will be to progress individual Principal Partner procurements.
- Seek to capture lessons learnt within recent procurements by Principal Partners within PRACE and complementary international procurements.

2 Overview of European Commercial Procurement Procedures

Deliverable D2.3.1 has made a good start in defining the various phases of the procurement of a system, identifying the roles and responsibilities of the various stakeholders and fleshing out the detail of some of the processes and procedures. Other tasks within work package 7 address the issues of developing risk registers, benchmarks, technical requirements.

The EU Procurement Directives (2004/18/EC of the European Parliament and of the Council 31st March 2004 on the coordination of procedures for the award of public works contracts, public supply contracts and public service contracts) set out the legal framework for public procurements. They apply when public authorities and utilities seek to acquire supplies, services or works and set out procedures which must be followed before awarding a contract when its value exceeds set thresholds, unless it qualifies for a specific exclusion – for example – on the grounds of national security.

For planned purchases above the threshold value a Prior Indicative Notice (PIN) may (optionally) be published in the OJEU Journal, typically at the beginning of a financial period. The PIN must contain as much as possible of the information normally published in a contract notice and must have been despatched to the Official Journal no less than 52 days and no more than 12 months before the date of despatch of the relevant contract notice. Where a PIN has been published in advance the timescales under the Open Procedure can be reduced from 52 days to 36 days and exceptionally 22 days and under the Restricted Procedure from 40 days to 26 days. The time limits run from the date on which the contract notice is sent in Open Procedures and the date on which the invitation to tender is sent in Restricted Procedures.

2.1 Open Procedure

Anyone can bid: suppliers respond to a notice in the Official Journal, all interested suppliers will be sent an Invitation to Tender (ITT). Unless a PIN (see above) has been issued the purchaser must set a response date of a minimum 52 calendar days from the original notice date. Most organisations use this procedure whenever possible as it opens opportunity to the widest competition. This procedure is clearly inappropriate for the procurement of very technical high risk contracts since there are in practice very few companies able to supply and the effort of reviewing tenders is large.

2.2 Restricted Procedure

The number of tenderers (suppliers) may be restricted to at least 5, if available, and only those suitable applicants (assessed by a business questionnaire) invited to bid may do so. A minimum 37 days must be given for expressions of interest. Shortlisted suppliers will be sent an Invitation To Tender and allowed a minimum 40 days to respond (unless a PIN has already been issued). Most organisations will use this procedure when there is a probability of high levels of interested vendors in the supply market which would result in difficulties and inefficiencies in the tendering process.

2.3 Competitive dialogue procedure

Following an OJEU Contract Notice and a selection process, the authority then enters into dialogue with potential bidders, to develop one or more suitable solutions for its requirements and on which chosen bidders will be invited to tender.

2.4 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.

2.5 Comments

Public Authorities typically have a free choice between the open and restricted procedures. The competitive dialogue procedure is usually only available where the contract cannot be awarded under open or restricted procedures. The negotiated procedure can only be used in very limited circumstances.

Under restricted, competitive dialogue and competitive negotiated procedures (those where a call for competition is required by advertising is advertised in the OJEU) there must be a sufficient number of participants to be selected to proceed to the tender stage to ensure genuine competition. The regulations require a minimum of five for the restricted procedure and three for competitive dialogue and negotiated procedures.

In the standard procurement process the stages involve:

- Specification requirements – using performance specifications rather than branded technical specifications
- Selection stage – the rejection or selection of candidates based on:
    - Unsuitability – convictions, bankruptcy, failure to pay taxes;
    - Economic and financial standing – financially sound;
    - Technical capacity and ability.
- Award stage – either on the basis of “lowest price” or various criteria for determining which is the “most economically advantageous tender (MEAT)” which is normally consistent with value for money.

There are restrictions on the use of post tender negotiation under the open and restricted procedures – there can be no negotiation on price.

2.6 Pre-commercial procurement

Pre-commercial procurement 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 EU has enormous purchasing power of order B1700 € pa of which less than B3 € is used for procuring the research and development of new products and services. The US spends four times as much in areas such as health and energy. By engaging Europe’s public procurers in

buying the development of new innovative products and services it is hoped that this will provide better public services and greater value for money and give European high-tech industry a chance to be the first to market. Some of the barriers to adopting pre-commercial procurement are the European procurers are fragmented, averse to risk and may have to overcome legal barriers to cooperate across borders.

Pre-commercial procurement typically occupies the product development pipeline starting with solution exploration, moving on to prototyping and then delivering a first limited volume of products/ services (a test series). The products then enter into the realm of commercial roll-out purchased through the standard public procurement procedures discussed earlier. This type of R&D procurement strategy has been used heavily in the US in particular by DARPA and the Department of Energy to support R&D for supercomputers, addressing public needs with long-term procurement plans and a competitive multi-supplier approach. Best practice is to procure in the steps discussed above to reduce risk, to involve multiple suppliers, to work across borders to develop larger markets and foster standards.

Access within these schemes is typically restricted to local industrial suppliers but has to be organised in compliance with WTO rules discussed in the: "AGREEMENT ON GOVERNMENT PROCUREMENT" at http://www.wto.org/english/docs_e/legal_e/gpr-94_e.pdf).

Legally it is the procurement of R&D services that falls under an exception of the WTO Government Procurement Agreement and the EU public procurement directives. For this exception to hold, the value of the pre-commercial procurement contract must consist of 50% R&D services. The value of any products procured in the contract cannot exceed the value of the R&D services covered by the contract. Under this condition risk-benefit sharing between the procurers and the suppliers and restriction of the tender to the European Internal Market is allowed. Eligibility criteria on suppliers related to the domestic location of the company office and/ or governance structure are not geared to creating growth hence to compete for a pre-commercial procurement contract a company should not necessarily have to be European or European owned as long as it is committed to locating the centre of gravity of the R&D and operational activities related to the contract in the EEA. The award of a contract through an open tendering procedure according to market conditions where the risk-benefit sharing conditions are fixed in advance of tender publication for each pre-commercial development phase then no state aid is assumed to be involved.

So in summary pre-competitive procurement could provide PRACE with an exciting opportunity to drive forward the development of innovative HPC systems to solve major societal challenges – the key issues are:

- 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.
- European capability or capacity to potentially meet this need.

Within the broader procurement strategy then the most appropriate route must be driven by the specification of requirements, whether these can be met cost-effectively by a broad market or whether there is a need to tweak current commercial offerings or invest significantly in R&D.

3 Overview of Procurement Objectives

Deliverable 2.3.1 defined a number of stages important in the procurement of systems:

- Justification and Elaboration – the science and business case
- Prototype Evaluation and Tender Design to inform the Specification of Requirements and Decision to engage in Tender – accepting the market's capacity to supply
- Investment Decision and Implementation of the contract
- Installation
- Acceptance and Pilot use
- Contract Closure
- Assessment – lessons learnt

It is interesting to take a step back and reflect on the objectives for the procurement of an individual system, for example those listed in Table 1 below:

Objective	Comment
Requirements	Are the requirements for: • a stable production platform that will deliver good performance for a general set of applications for a broad range of users with mixed HPC experience; • a specific architecture that will deliver very high cost effective performance for a limited range of applications with an experienced user base, or, • to inform innovation of new designs for future generations of systems? This encompasses a spectrum of applications performance, availability and ease-of-use requirements.
Flexibility of procurement	What are we trying to optimise here – is it the initial system, is it the evolution of the system and/ or is it the integrated performance of the system over a period of time?
Lifetime costs	Important to consider total lifetime costs in evaluating value for money. These should include not only the cost of the system per se but the maintenance, system support, infrastructure (eg special floor/ cooling/ uninterrupted power supply/ air and water quality), running costs in particular electricity and system support costs.
Procurement Process	Which national/ international procedure and why? Time/ flexibility/ cost of procurement/ non-standard requirements?
Minimise risk of supplier default	Assurance of sustainability/ ongoing competitiveness of supplier with respect to general HPC market and with respect to the specific regional procurement through pre qualification questionnaires or warranties and guarantees.
Maximise value for money/ most economically advantageous tender	Benchmarks and availability acceptance tests. Risk transfer and added value.

Table 1 : objectives of procurement for an individual system

In the following section we review procedures and processes followed in recent procurements and draw some conclusions to inform PRACE procurements.

4 Overview of Recent PRACE Principal Partner and International Procurements

In the following Table 2 we summarise recent procurements by the Principal Partners and a number of international procurements. Details can be found in the Annex. These include:

- 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

We have approached a number of other overseas sites that have recently procured systems but there has been a reluctance to provide comments because of issues to do with commercial confidence and the availability of non-english documentation.

At this point the various organisations that have contributed support a reasonable spectrum of requirements which have resulted in different procurement strategies which are summarised below. We follow this with some general conclusions.

Principal Partners	Requirements	Flexibility	Lifetime Costs	Procurement Process	Contract/ negotiation issues/ Supplier default	Benchmarks/ performance/ availability/ acceptance	Evaluation
NCF	Capability computing National supercomputer (broad range of applications) Minimum core performance/ number of cores per node, memory per core and memory per node (heterogeneous), tough I/O and high. Quote for 40/ 60/ 80 TF systems.	Midlife upgrade	Evaluation of total cost of ownership.	16 month process. Input to RfP via vendor briefings. Followed by Open OJEU tender. First step open to any vendor evaluated by response to requirements – score of 1-10. 3-5 vendors then invited to benchmarking and on-site visit (5-days).	Disclosed lower bound on budget hence request for prices of a range of systems with different performance levels.	Normal load plus future extrapolations. 12 applications – 16-128 processors. High memory and I/O requirements. As is and stand-alone, optimised plus throughput test on 256 cores and optimised. Opportunity to submit revised benchmarks.	200 requirements – Mandatory and desirables given prefixed weightings (1-10) judged by experts. Categorized in 7 groups including finance and future technology. Final evaluation using weighting of score of 1-10 for desirables, benchmarking and site visit Vendors aware of ranges of weights.
EPSRC	Capability systems – typically two services run for 6 years overlapping for 3 years. Minimum requirements specified around memory and performance.	Performance increase after two years. Option for further refresh.	All costs included in evaluation.	Recent service procured as independent components for science support, hardware and hosting. Restricted Procedure – specification must be clear. Bidder briefing, Pre Qualification Questionnaire. Outline Proposals, shortlisting, Invitation to Negotiate, Best and Final Offer.	Supplier default evaluated through pre-qualification questionnaire.	Mixture of kernels, applications and throughput tests. Availability test for 10 days.	Bids reviewed against multiple assessment criteria.

Principal Partners	Requirements	Flexibility	Lifetime Costs	Procurement Process	Contract/ negotiation issues/ Supplier default	Benchmarks/ performance/ availability/ acceptance	Evaluation
France	General purpose HPC system (SMP cluster architecture) with high disk and bandwidth requirements.	Options for extending the equipment and add novel architecture capability.	Acquisition and maintenance/ support costs for four years.	Competitive dialogue which can iterate the specification. Most recent procurement was for the collaborative design of a unique system with high communications and I/O – option to purchase system.	Budget disclosed after checking that all bids exceed budget. Revised bids evaluated on disclosed budget.	Mix of low-level and applications benchmarks with a particular emphasis on I/O performance.	300 requirements much reduced for collaborative design project to allow flexibility.
Juelich	Procures both capacity and capability systems.			Capacity systems procured through competitive dialogue for general purpose systems, an R&D contract to explore emerging technologies and a negotiated procedure for capability systems.	Vendors cannot commit to performance levels on leading edge capability systems. Full configurations not available so acceptance linked to performance on smaller partitions on full system.	Low level and applications for general purpose systems. Linpack and comms test.	Optimal solution to fixed specification for given budget for general purpose system. R*D contract allows influence on specification and design.
BADW-LRZ	Large systems are pure capability. Migration system needed in advance. Installation normally in two phases.	Phase 2 installation technology left open as long as possible.	Investment plus energy costs (direct plus cooling overhead).	Large systems procured through a competitive dialogue following an intensive phase of technical reviews with all possible vendors.	Contracts care for delays and technical problems. Penalties for failed benchmark commitments.	A mix of kernel and application programs designed to simulate the expected load.	Exclusion criteria for inappropriate offers. Best performance commitments for fixed price. Assessment of usability needs.
BSC	Cluster of 2560 JS21 blades			Cooperation with IBM emanating from a joint research institute which designed a system.			

Principal Partners	Requirements	Flexibility	Lifetime Costs	Procurement Process	Contract/ negotiation issues/ Supplier default	Benchmarks/ performance/ availability/ acceptance	Evaluation
NERSC	10% Capability plus more general user service	Options to purchase additional memory or CPUS.		18 month process form publication of RfP to start of service.		System component tests, kernels, reduced applications, full applications and composite tests. Asis and optimised. Demanding acceptance tests.	
ECMWF	Very high-availability system required for limited number of applications. Performance requirements focus on minimising time to solution for various job run mixes. Specified need for pre and post-processing and test and development systems.	Mainly on timing of various anticipated performance phases.	Included infrastructure and electricity cost-calculator for total lifetime costs.	Restricted procedure. 9 months form open briefing to contract signature for a system to be delivered 9 months later.	Specified lower and upper budgets. Penalties spelt out for missing preferred milestone dates.	Functional and performance tests on applications, throughput and disk access and network kernels. As is and optimised results with clear limits on optimisations. Strong emphasis on resilience.	Over 300 requirements including mandatory, highly desirable, desirable and requests for information. Best performance level and profile of performance.

Table 2 : summary of Procurements

Sections below contain general comments drawing on the table and the Annex 6.

4.1 Requirements

In preparing the requirements there should be extensive consultation with the vendors a year ahead of time on systems that may be available over the following 5 years.

Requirements are typically categorized into groups including technical, financial and future technology roadmaps.

The requirements specifications describe the need for a range of systems to meet the needs of a broad range of users undertaking a spread of science with many hundreds to many thousands of users. A subset of these users may have capability applications the need to accommodate a broad base of users and which emphasises sustained performance, systems integration and high availability which tends to result in a conservative procurement strategy – purchasing commercially available systems form large companies with appropriate support infrastructures.

It is sensible to include a test system in the requirements enabling software to be evaluated in a realistic environment before being rolled out onto the production platform. It is desirable that the test system is delivered before the main system as this enables the development of systems management procedures and initial optimisation of applications.

It is sensible to specify a front-end system that can be used for pre and post-processing of output. This usually implies that a fraction of the system has high memory nodes with enhanced I/O to the filestore.

Deliverable 7.5.1 presents best practice on the specification of 'Technical Requirements for Petaflop system(s) in 2009/10' derived from recent procurements.

Accurate specification of site requirements is increasingly important with a clear understanding of peak/ sustained power, footprint and weight, cooling requirements and detailed airflow/ fluid coolant modelling, particulates specification and fire suppressant requirements. Site issues are being developed further in Task 7.3.

4.2 Flexibility

A number of the recent procurement highlighted flexibility in requirements provision through the timescale of the procurement. This included:

- Specifying a profile of increasing requirements as a function of time with key system parameters such as compute performance, memory, disk, backup and I/O performance.
- Options for purchase of reduced/ increased system configuration, increased benchmark performance, filesystem size and performance, memory, interfaces, future technology evaluation systems, additional on-site support.
- Allowing slippage in delivery for some compensation in system performance.
- Accommodating technological changes that may be proposed to save money, to improve performance or to save energy ot to accommodate increased capability/ capacity needs.

4.3 Procurement Process

The procurement processes range from joint R&D exercises to standard restricted or dialogue procedures to meet the differing service objectives.

4.4 Benchmarking

Most procurements use very similar benchmarking procedures encompassing component tests, kernels, applications, throughput and availability tests. These are being addressed in various other deliverables within this work package, in particular, Deliverable 7.5.1 - Technical Requirements for Petaflop/s system(s) in 2009/10 and Deliverable 6.3.2 – Final Benchmark Suite.

4.5 Acceptance

It is normal to require that the system is assembled and tested at the site - including all hardware installation and assemble, burn in of all components, installation of software and approved production environment, tests and benchmarks addressing functionality, performance, reliability and quality and run benchmarks to demonstrate performance commitments. Systems tests should demonstrate a reasonable amount of time taken to boot the entire system from a cold start to the production state and the reverse power off option.

In addition it is desirable to run a Factory Test prior to shipping covering power up and down, UNIX commands, monitoring software, reboot functions, power cycle from console, configuration testing benchmarks and a variability test expressed as some meeting some maximum % variance over a number of days. It may not be possible for some vendors to test the full system because of limitations on the factory's power and cooling.

Availability tests usually typically run for 10-30 contiguous days in a sliding window of 20-60 days and require typically 98% availability – often running a selected user service. Failures include unavailability of nodes, inability to access the file system, inability to login, unavailability of full switch bandwidth, inability to launch batch submission.

Functionality demonstrations are run on the configuration that will go into production and include remote monitoring, power control and boot capabilities, network connectivity, file-system functionality, batch system, system management software program development environment and UNIX functionality.

Systems tests should include time from a cold start to production state in a reasonable time period and single node power tests.

Performance tests are structured around benchmark suites ranging from simple component tests through to applications performance to throughput tests.

4.6 Evaluation

Evaluation procedures for the procurement of standard systems is fairly uniform in concept with desirable requirements given weights which are summed to give a weight for a category of requirements. These are then input into a weighted sum incorporating commercial factors concerning the tenderer, the extent of risk/ reward sharing and benchmark results. The weightings clearly need to be balanced against the objectives of the procurement.

4.7 Contract

Typically a fixed term contract with line items for a test system, main computer systems, maintenance and servicing, and performance metrics as a function of time and other negotiated features and deliverables. Most contracts have a Change Control Procedure, named personnel, insurance, and priced options.

The contract needs a clear specification of the site preparation, subcontractor commitments warranties and representations, transferable software licenses, access to facilities and network policies and procedures.

Where performance metrics are mandatory it may be necessary to include charges for additional infrastructure, system support and running costs should the tendered configuration need to be increased.

Need for ASIS/ optimised / throughput benchmarking and tests for general user systems in the production systems environment.

5 Summary and Conclusions

Most of the recent procurements have followed very similar routes spelt out perhaps unsurprisingly in Deliverable 2.3.1. Much of the best practice and lessons learned from recent procurements are and will be captured within other Deliverables within Workpackage 7 for example:

- Deliverable 7.3 – Report on Installation Requirements and Availability at European sites.
- Deliverable 7.5.1 - Technical Requirements for Petaflop/s system(s) in 2009/10
- Deliverable 7.6.2 – Pre-qualification questionnaire - used to assess the current financial standing and medium to long term financial viability of bidders, their capability to produce a viable technical solution and their propriety.
- Deliverable 7.6.3 – Evaluation criteria and acceptance tests – building on the earlier work of Deliverable 2.3.1 to develop appropriate metrics tailored to the requirements of the procurement.

Whilst the emphasis within this deliverable has been on the procurement of a single system there is a number of broader strategic issues that PRACE still needs to resolve before developing a coherent procurement strategy:

- What is PRACE procuring – is it a service (system, site, scientific support) or the individual components – what are the timescales for these procurements and should they be different for the different components?
- Will PRACE procure a single system at a time or are the individual procurements part of a portfolio of procurements meeting a range of requirements including the need for general production systems, the need for specific systems for specific applications and/ or the need to innovate in informing the next generation of systems design (over and above the activities within WP8). Clearly these considerations depend on the detailed science and business case which would translate into very different requirements and choice of procurement procedure?
- Will PRACE conform to the cycle or operator model – this could effect the legal form undertaking the procurement and the legal procedures open to then to conduct the procurement? This may effect the cost of systems (tax implications) and the ability to attract industrial activity (not to exceed 10% for example to avoid VAT in the UK).

This deliverable highlights important issues that the Management Board may need to take into account.

6 Annex

In the following we present summary reports on the procurement of HPC systems by the Principal Partners.

6.1 Procurement by NCF, Netherlands

Peter Michielse, NCF, November 14, 2008.

Summary
This document described the recent procurement of capability computing equipment by the Netherlands National Computing Facilities Foundation (NCF), for a new national supercomputer to be installed, hosted and run by SARA Computing and Networking Services in Amsterdam. The new national supercomputer is targeted to conduct public research by scientists at the Dutch academic institutes (universities and research institutes). The procedure has taken about 16 months, starting with gathering input for a Request-for-Proposals (RfP) at SC05 in Seattle, actually writing the RfP and starting an open European tender procedure in the first half of 2006, accompanied by extensive application benchmarking and site visits at the selected vendors in late summer and fall 2006, making a decision on which vendor to finalise a contract of delivery, which has been signed in early spring 2007.

Introduction
NCF is the Dutch funding agency for the national computer infrastructure for scientific research. This means actually selection and purchase of high-end computing and data storage equipment, which encompass the regular budgets of Dutch universities and research institutes for such scientific equipment. With respect to computing equipment, the policy of NCF is to invest in both capability computing and capacity computing. This procurement has been targeted to capability computing, which has been reflected in the RfP. NCF has a long tradition of supercomputing equipment: this has started in 1983 with CDC Cyber 205 systems towards Cray Y-MP and C90 vector systems, with a transition (in 2000) to massively parallel shared memory SGI Origin and Altix systems. This procurement is for the replacement of the 3.2 Tflop/s SGI Origin3800/Altix3700 combination, with a total of 1440 processor cores. The supercomputer systems will be hosted and managed by SARA in Amsterdam. SARA becomes owner of the system, while NCF (belonging to the Dutch organization for scientific research NWO) as the funding agency controls access (peer review) and usage of the system. It also means that SARA and NCF have intensively worked together in this procurement.

Overall schedule of procurement process

Information:	November 2005 to December 2005;
Preparing RfP:	January 2006 to April 2006;
Preparing benchmark:	January 2006 to April 2006;
European tender procedure:	May 2006 to June 2006;
Judgement of proposals:	July 1006 to August 2006;
Benchmarking:	May 2006 to September 2006;
Reporting:	October 2006;
Decision:	November 2006;
Finalisation of contract with selected vendor:	December 2006 to March 2007.

Information
For many years, NCF has carried out research on supercomputer architectures, processors and interconnects, and it publishes these on an annual basis [1]. Furthermore, NCF conducts

research with respect to academic supercomputing in Europe, to remain aware of developments in, and installations of, high-end computing equipment across Europe. The research is published on an annual basis [2].

Through this research, NCF started to gather information on available systems at SC05 in Seattle. This has been done very carefully by (non-disclosure) vendor meetings. All relevant vendors were informed, appointments for two-hour meetings were scheduled and were executed at SC05. NCF prepared a two-page questionnaire with subjects to be addressed by the vendor during the meeting, so that vendors have gotten the opportunity to prepare themselves. The subjects covered varied from architecture details to processor details and roadmaps to financial information and maintenance/services profiles.

After SC05, NCF has collected all the information, and has asked each vendor to verify its own information, as presented to them by NCF. After verification, NCF has used this information for building up the requirements in the RfP.

Preparing the RfP
Before writing the RfP, NCF had concluded that the user requirement was for a capability system. This has led to a set of requirements, for which the most important were:

- Single-core performance at least 5 Gflop/s;
- At least 8 cores per shared memory compute node;
- At least 4 GB memory per core, at least 64 GB per compute node, for 80% of the nodes;
- At least 8 GB memory per core, at least 128 GB per compute node, for 20% of the cores;
- Quite tough IO requirements.

On request, we can make available the RfP document.

Overall, there were around 200 requirements, some of them mandatory (knock-out), other valued by a weight with respect to importance. In practice, this turned out to a lot (maybe too much). Each requirement had its own weight. Vendors were aware of the range a weight was in. The actual weights have been fixed by NCF before receiving the proposals.

Further, the requirements in the RfP have been categorized in groups (about 7 groups). There were technical groups, but also a financial group and a group on future technology and roadmap. Each group got a certain weight again. Each vendor had to pass the mandatory requirements, after which the vendor answers/proposal to the other requirements were judged by a team of experts. This has led to a number between 1 and 10, the highest number being the best qualification.

An additional remark on available budget. NCF has not disclosed the available budget in the RfP, since we ourselves only knew the lower bound. Instead of that, we have asked the vendors to make a financial quote for 40 Tflop/s, 60 Tflop/s and 80 Tflop/s systems.

Preparing benchmark
All applications in the benchmark suite have been recruited from the normal load on the Dutch national supercomputer, with the request to the developers for preparing their code for a large amount of processors, with relevant input sets (i.e. for execution on a future, much larger and faster system). This has led to around 12 applications, running on 16 to 128 processor cores, with high memory and IO requirements. Each application had to be run as

supplied and standalone, optimized performance as well, and also as part of a throughput test on 256 cores.

Each interested vendor has been given the opportunity to download the full benchmark environment at the beginning of the tender procedure, in early May 2006. Reporting the benchmark results was not required in the first phase, but only during the site visits to the selected vendors. So each vendor took some risk in starting working the benchmark, as it may not be selected on the short list.

European tender procedure
Actually, the tender process was split into two steps. The first step was open to any vendor, and was completed by an evaluation of the requirements. This has led to the number between 1 and 10, as noted earlier. Then, a shortlist of 3 to 5 vendors was made by the NCF board to go into the second step. This second step consists of extensive application benchmarking on-site at the vendor's premises. The benchmarking results give rise to a second number between 1 and 10, and the actual quality of the visit also to a number between 1 and 10.

The NCF board was advised by a selection committee. This selection committee has been set up by the board of NCF, consisting of representatives of the sciences using the HPC equipment. This selection committee has been involved in the RfP as well. The selection committee was advised by NCF and SARA staff.

Assessment of proposals
As has been mentioned before, this has been done based on the non-mandatory requirements, their weights, and the weight of their group. This has led to a single number between 1 and 10. Based on this, the selection committee advised the NCF board on the vendors that would proceed into the second step.

Benchmarking
Vendors have gotten the opportunity to start working on the full benchmark in May 2006. NCF and SARA staff (3-4 persons) has been visiting 3 vendors for actual running of the benchmark set (which was/should have been prepared by the vendors). Each vendor has been visited during 5 full days in September 2006. Each vendor had the opportunity to send in improved benchmark results up to September 30, 2006. Apart from benchmarking, each vendor received in advance a set of questions from NCF, based on their proposal and in most cases asking for clarification by the vendor.

Reporting
Early October 2006, NCF has collected all benchmark results and has prepared a report, containing the results of the requirements, the benchmark and the site visit itself. This report has been discussed with the selection committee. The selection committee, based on the reporting, has made its advice to the NCF board., which has taken the decision to start finalizing the contract with the selected vendor.

Basically, the report contained three numbers:

- One for the proposal;
- One for the benchmark results;
- One for the site visit.

NCF has calculated a weighted average of these three numbers (weight were known to the vendors with respect to the range they were in). The selection committee took the vendor with the highest number and has advised to the NCF board to start finalizing a contract with this vendor.

Decision
Based on the advice of the selection committee, the NCF board has decided in November 2006 on the vendor to try to finalise a contract with. At that point in time, the available budget for the investment of hardware and the budget for the full exploitation was available to NCF.

Finalisation of contract with selected vendor
This has not been an easy task, and has taken quite some time, but finally a contract could be signed.

References
[1] Overview of recent supercomputers, A.J. van der Steen, NCF, annual publication;
[2] Academic Supercomputing in Europe, R. Llurba, NCF, annual publication.

6.2 Recent procurements in French National Capability Computing Centres

Contribution by J.-Ph. Nominé, CEA, with the help of François Robin, GENCI, Stéphane Requena, GENCI, and Christine Ménaché, CEA/CCRT.

This section does not intend to describe all the details of the recent procurements for French national equipments, but rather to highlight salient features of the different approaches.
All procurements used quite standard technical approaches with benchmarks, accurate lists of requirements categorized into groups, evaluation criteria etc. So we won't insist on these aspects unless we want to emphasize some specificity.

French national supercomputing centres

CINES (**Centre Informatique National de l'Enseignement Supérieur,** http://www.cines.fr/) is a national computing facility depending on MESR (Ministère de l'Enseignement Supérieur et de la Recherche, French Ministry of Higher Education and Research). Located at Montpellier, it provides computing resources and support to the community of higher education and public research.

IDRIS (**Institut du développement et des ressources en informatique Scientifique**, http://www.idris.fr/) is the main CNRS (Centre national de la recherche scientifique) computing centre. Located at Orsay near Paris, it provides computing resources and support to the scientific and public research community.

CEA operates two distinct computing facilities at Bruyères-le-Châtel site near Paris: **TERA** for defence applications (classified facility) and **CCRT** for research and industrial applications.

CCRT (**Centre de Calcul Recherche et Technologie,** http://www-ccrt.cea.fr/) has the particularity of having also industrial partners who are real shareholders, not only stakeholders (they co-invest in the equipment). The same CEA division operates both TERA and CCRT facilities, which fosters a lot of synergy from procurement until operation through

installation and deployment, even if the two supercomputing centres are settled in two separate buildings and have different security policies.

In 2007 GENCI, a national agency, was created in order to coordinate the French HPC policy for research. GENCI coordinates new national research equipment investments and usage at CNRS/IDRIS, CINES and CEA/CCRT, the three national tier-1 centres (CEA/TERA remains under the separate authority of the French Ministry of Defence). The first two new procurements coordinated by GENCI in 2008 were CINES equipment renewal and CCRT upgrade we describe below.

In 2008 the National Jacques-Louis Lions Centre was created as an encapsulation of IDRIS and CCRT which are actually in the same administrative district ("département de l'Essonne" - 91) and close to each other (20 km). The objective was to strengthen their coordination and visibility and to propose to GENCI a unique entity for the French candidature for hosting a PRACE Petaflops class infrastructure. But the two supercomputing centres remain separate physical entities belonging to their respective CNRS and CEA establishments.

Recent procurements

Table 3 below summarizes recent procurements in the aforementioned computing centres, which we then briefly comment.

	TERA10	**CCRT-B**	**CINES**	**TERA100**
Organization	CEA	CEA & GENCI/CEA	GENCI/CINES	CEA
Site	Bruyères (classified)	Bruyères	Montpellier	Bruyères (classified)
Equipment	60 TF Cluster	50+100 TF Cluster + 200 TF SP GPU extension	150 TF Cluster	1 PF Cluster
Vendor	Bull	Bull	SGI	Bull
Procurement dates	2004 Signed December 2004	11/2005-09/2006 Signed December 2006	10/2007-04-2008 Signed April 2008	01/2008-06/2008 Signed July 2008
Installation	12/2005	April 2007 Extensions: April 2009	July 2008-12-01	2010
Procedure	*Appel d'offre sur performance* *	*Appel d'offre sur performance* *	Competitive dialog	Competitive dialog + R&D contract

* this type of procedure, literally "tender on performance", was later replaced by the "competitive dialog"; it freezes the specifications and iteratively optimizes the solution, whereas competitive dialog can really iterate on the specifications.

Table 3 : examples of recent procurements at French national centres

TERA10

TERA 10 is a general-purpose cluster of 60+ TF based on Intel Itanium Montecito dual-core CPUs and 16 core compute nodes and uses a Quadrics QSNet III (Elan4) interconnect.
The procedure used was a "tender on performance", with accurately defined specifications and benchmarks but allowed iterative optimization of the proposed solution under the given budget. Almost 300 specification items were used for TERA10 procurement.

CCRT-B

The so called CCRT-B procurement is a major upgrade of CCRT equipment endeavoured at the end of 2006. The procurement phase spanned most of 2006; the final machine was installed in April 2007 and was in full "open" production in September 2007 after an intensive "grand challenges" period during 2007 summertime.
The machine is a general-purpose cluster of 50 TF based on Intel Itanium Montecito dual-core CPUs and 8 cores compute nodes and uses an Infiniband interconnect.
"Tender on performance" procedure was also used, and the project benefited from CEA experience with TER10, concerning requirements, benchmarking, installation and deployment. This was before GENCI was created.

The contract has several options which allow for various kinds of extensions. In 2008, some of these options were used to order extensions of the equipment, funded by GENCI. A specific protocol was signed with GENCI to ensure a well-defined coordination within the scope to the existing contract. These extensions are:

- A general purpose cluster part of ~100 TF, consisting in 1068 Bull Novascale R422 nodes based on Intel Nehalem quad core processors
- A more specialized part based on 48 new generation NVIDIA S1070 GPU servers, for a single precision peak performance of 192 TF

CINES
The new CINES machine was funded and procured by GENCI in close collaboration with CINES and used a competitive dialog procedure. It is a 147 TF peak performance Alitx ICE 8200 cluster of bi-socket nodes with quad-core Intel Harpertown processors, with an Infiniband interconnect network. Launched in October 2007, the procedure followed the conventional stages of a competitive dialog. The vendors were selected in December 2007, the competitive dialog by itself lasted until February 2008. The final choice was made in March 2008. Shipping and installation took place between May and July. July to September 2008 were dedicated to Grand Challenge applications before full production deployment.

TERA100
TERA 100 will replace TERA10 in 2010. It will be a 1 PF peak performance cluster. The procurement stage lasted from January to June 2008 and a contract was signed in July 2008. This stage consisted in selecting a partner for collaborative design of the system. The spirit of the approach was indeed quite specific, so as to handle the open and challenging design of a capability, general-purpose production cluster with very high demands on interconnect and I/O bandwidths. The contract mostly deals with a 2 years R&D stage, where the machine acquisition is a final option. By contrast with previous CEA specifications, the list or requested criteria is a quite condensed one, focusing on main sizing parameters or critical features, so as to allow the highest possible design flexibility. This approach, together with the sharing of some intellectual property, fosters a very tight collaboration between the vendor's and the customer's teams, far beyond the usual competitive dialog method. For CEA it matches the need for designing a unique object suited to its specific production.

Conclusion
Those last 3 or 4 years, the French national supercomputing centres have mostly used a conventional procurement approach when established or mature technologies are suited to the needs – competitive dialog or equivalent previous type of procedure. After a solid foundation has been given to these procurements, using options defined in the initial contract is a good way to offer some flexibility for upgrades along the usual 3 to 5 years equipment lifecycle.

The more we go into challenging design for petascale machines, the more it is important to be able to iteratively design unique objects whose exact specifications can only iteratively stem from close and long term dialogs between computing centres and vendors, like in pre-competitive procurements. Not formally being a pre-competitive procurement, TERA100 contract however illustrates an approach with that kind of enhanced flexibility.

6.3 Procurement by Jülich, Germany

M. Stephan, FZJ

Introduction
The Jülich Supercomputing Centre (JCS) is part of the Research Centre Jülich (FZJ) and one of the three members of the Gauss Centre for Supercomputing (GCS) beside the High Performance Computing Center Stuttgart (HLRS) and the Leibniz Supercomputing Centre (LRZ).

The Gauss Centre for Supercomputing (GCS) is the alliance of the three German national supercomputing centres providing world-class supercomputing power for computational science and engineering for Germany and Europe. The major funding for supercomputer infrastructure is given by the Federal Ministry of Education and Research (BMBF) and the state ministries for research of Baden-Württemberg, Bayern and Nordrhein-Westfalen. But also research organisations contribute to the budget.

JSC pursues a dual system approach with a High-End, general purpose system for capacity computing and a more special Leadership system for high-scaling, capability computing to satisfy different user needs. JSC has a long tradition in hosting supercomputer systems starting 1983 with Cray X-MP, Cray Y-MP, Cray T90 and Cray T3E, Intel Paragon and since 2003 an IBM p690 cluster, Blue Gene/L and Blue Gene/P.

Recent procurements
Table 4 summarizes recent procurements at FZJ, further commented in the section below.

	JUMP	JUROPA	JUGENE
Equipment	5.8 TF Cluster	200+100 TF Cluster	223 TF MPP system
Vendor	IBM	Bull/SUN	IBM
Procurement dates	08/2002		06/2007
Installation	07/2003 Extension 03/2004	04/2009	10/2008
Procedure	Competitive dialogue procedure	R&D contract	Negotiated procedure

Table 4 : examples of recent procurements at FZJ

JUMP
JUMP was a 5.8 TF general-purpose cluster based on IBM Power4+ processors (p690 systems) connected with an IBM High Performance Switch. This procurement used a competitive dialogue procedure with a benchmark suite (low level and scientific applications) and defined specifications to develop and find an optimal solution under the given budget.

JuRoPA
JuRoPA (Juelich Research on Petaflop Architectures) is a research project of FZJ together with Bull, Intel, ParTec and Sun to develop a general purpose cluster to investigate emerging cluster technologies and achieve a new class of cost-efficient supercomputers. For FZJ this is a new approach to get more influence in design and specification of a new system.

JUGENE
Procuring a new system is always a challenge, especially when trying to take a huge step forward. JUGENE was the second Blue Gene/P installation world wide and the first of this size. Choosing benchmarks in such a case is very difficult because no vendor has a system of that size in his labs where benchmarks can be run. So beside Linpack we only have one benchmark (communication network test) that is proven to scale on such a huge number of nodes. With all normal application benchmarks there is the problem that a prediction of the scaling is very difficult. So no vendor is willing to take the financial risk of binding the acceptance of the whole system to an application benchmark run on the whole system. So for all application benchmarks we allow runs on (sub)partitions of the final big system. Those runs had to perform as measured on an existing smaller system.

With this system FZJ faced the first time, that the restrictions in the power and cooling capacity of the building can limit the range of systems that meet those requirements.

Conclusion
There is a high demand of different procurement approaches in the field of supercomputing. When looking for general purpose systems with standard technology and of existing size a competitive dialog or equivalent type of procedure seems to be the best approach giving enough flexibility to find a good over all solution.

Procurements of leading edge technology or of really big systems like the planed tier-0 systems in PRACE needs a high flexibility in the procedure to define and discuss all the challenges in such a project with the vendors and maybe adapt the specifications. It is necessary to have a fine grained planning how to react if given specifications can't be achieved. We can't expect that any vendor will take the financial risk of retracting a 100M€ system – particularly without the chance of selling the hardware to other customers – just by missing one or two specifications.

6.4 Procurement by Barcelona Supercomputing Centre, Spain

S Girona, BSC

Barcelona Supercompting Center – Centro Nacional de Supercomputación (BSC-CNS) is the Spanish National Supercomputing facility, and is a collaboration between MICINN (Ministerio de Ciencia e Innovación), DIUE (Departament d'Innovació, Universitats i Empresa), and UPC (Technical University of Catalunya), with the objective of performing research in Computer Sciences, Life Sciences, Earth Sciences and Computer Applications in Science and Engineerng, and providing Supercomputing Services to the Scientific Community.

The main HPC system operated by BSC-CNS is MareNostrum. A system provided by IBM and consisting of a cluster of 2560 JS21 blade servers providing 10,240 IBM Power PC 970MP processors at 2.3 GHz for a peak performance of 94 Tflops.

MareNostrum

The acquisition of MareNostrum provides an interesting case-study for future PRACE procurements as it was based on a long term collaboration between BSC and IBM for research and development.

Cooperation started in 2000 with a joint research institute where a list of research topics were defined and developed together, in order to create a system designed for the needs of the Spanish researchers. Some of the research topics included performance analysis tools, networking, programming models, mathematical libraries and file systems. BSC-CNS identified the applications which were to be run on the new system and the approach was to jointly prepare the required technology. The system was provided in 2004 and was ready for operation from the very first day with people skilled in operations and applications running. The first system was composed of JS20 blades, single core on dual chip blades, and extended to the current configuration in 2006.

During the installation and system installation, proper software adjustments and requirements were included for the 2006 upgrade, to support some scalability concepts. Joint research and cooperation was required to perform those objectives.

The long term collaboration and knowledge sharing between BSC-CNS and IBM teams has benefits for both sides of the relationship. BSC-CNS requirements, particularly for cutting edge HPC systems, are more likely to be understood correctly and IBM have the opportunity to advance their own research and development. BSC-CNS benefit from resources donated by IBM and the end result is an HPC system that meets the needs of the BSC-CNS users. This process has similarities with the GENCI TERA100 procurement and the pre- commercial procurement model.

6.5 Procurement by EPSRC, UK

J. Nicolson, EPSRC

HECToR Procurement Strategy Overview
EPSRC has traditionally procured high performance computing capability as a total service (i.e. facility's management, computational science and engineering support and hardware) as opposed to a hardware centric procurement model.

HECToR's predecessor, HPCx, was procured under competitive tender and awarded as a single contract to the UoE HPCx consortium for the full service length (six years). The contract made provision for two additional technology refreshes for HPCx restricted to the hardware partner in the UoE HPCx consortium (IBM). This provided a three-phase hardware refresh path with no change to the underlying service infrastructure in order to provide continuity of support.

The HECToR service was procured as three distinct components. To increase value for money and choice n the procurement, each aspect of the HECToR service was procured separately. This resulted in the following service topography:

CSE support was awarded for the full six years of the service to provide continuity of support.

Hardware provision was based on a "four plus two" model for the full six years of the service. As such, the contract awarded to the successful bidder ensured provision of the initial hardware solution and first technology refresh. An option is available for EPSRC to fund a second technology refresh.

Accommodation and facilities management, collectively referred to as Service Provision, was awarded after the hardware provider had been selected. This allowed the best option to be selected for the hardware and for the cost for modifications to existing facilities to be more accurately determined by prospective bidders. This ensured that any physical restrictions of the facility would not influence the final choice of hardware solution.

The result of this strategy is that there is no over arching contract for all three components of the HECToR service. The hardware is provided through a contract with the service provider who hold a contract for hardware and service with EPSRC, with EPSRC holding a separate contract with the CSE contractor.

The procurements followed standard EU guidelines and were managed under Prince 2 project management guidelines.

Procurement Procedure

Hardware
Due to the high technical content of the hardware solution and the degree of risk associated with this aspect of the procurement the Research Councils chose to follow a Restricted Procedure for procurement. Once the OJEU notice had been placed, the following processes were put into action the completion of each marking a procurement milestone.

- Bidder Briefing Day – Supercomputing 2005, public statement of the intention to procure a system and the broad technical scope and budget followed by 2 days of "one to one" meetings with prospective bidders.
- Pre Qualification Questionnaire – Used to assess the current financial standing and medium to long term financial viability of bidders, their capability to produce a viable technical solution and their propriety. This was followed by the Invitation to Submit Outline Proposal (ISOP) to those passing the PQQ stage.
- Short list and Invitation To Negotiate - The short list of bidders for the hardware was compiled following the assessment of the ISOP responses and consideration of the initial benchmark results. The Research Councils decided to restrict the shortlist to a maximum of five bidders at this stage.
- Following this a series of in depth meetings were conducted between the vendors and Research Council in order to refine their bids to produce a Best and Final Offer (BaFO) from each Vendor including all finalised benchmark results for both user and HPCC codes. The final bids were reviewed, scored against multiple assessment criteria, and ranked against each other.

CSE and Service Provision
Procurement in this instance was carried out utilising a Restricted Procedure. Similar in many respects to the above process, organisations wishing to bid for these aspects of the service first submitted details of their financial and technical capability for consideration by the Research Councils. From these expressions of interest, the Research Councils then selected a list of possible bidders to be invited to submit bids based on their technical and financial assessments only. A short list could then be generated, as opposed to requiring a large number

to vendors to submit detailed bids. This procedure is the most commonly used for services in the public sector and is suited to services which have high levels of technical or specialist requirements or significant financial strength in the supplier. This procedure, more importantly, demands that the services to be acquired are capable of clear definition and are well understood in the market. Consequently, the competitive dialogue stages are usually unnecessary (or for clarification) due to the defined nature of the requirements and equally definable nature of how those requirements can be delivered.

Best Practice Framework
Throughout the procurement project for HECToR each stage was subjected to a full external Gateway reviews to ensure that the strategy, process and business aspects of the procurement and service implementation were robust and fit for purpose. This was necessary in order to lend confidence to the Ministerial release of funding decisions necessary for projects of this scale and in order to provide external project assurance and where necessary guidance.

More information on the Office of Government Commerce Gateway Review processes can be found at:: http://www.ogc.gov.uk/what_is_ogc_gateway_review.asp

6.6 Procurement by ECMWF, UK

R.J. Blake – review of Invitation to Tender

Introduction
ECMWF runs a multinational medium range weather forecasting service that demands high availability, high performance for a few very high-profile applications with broader usage in a set of research activities. They are currently in the final throes of procuring a High Performance Computing Facility starting at the end of October 2008 to take over from the current service that is scheduled to end in March 2009. It will provide a production facility for about 4 years. The new system should realise a sustained performance of 20 TF when running the main application. An approved funding level (£42M) and a target funded level (£50M) minus the ancillary costs were detailed. Modular ancillary costs such as ECMWF's electricity costs and enhancements to the infrastructure would be used to evaluate a total cost of ownership.

The procurement objective was to acquire a HPCF with the best possible performance level and profile (progressive growth) for the funding available, taking into account the level of support to be provided, the need for reliability and resilience in the tendered solution and the technical requirements specified in the ITT. Proposals were invited for both the approved and target funding scenarios. ECMWF typically operates two identical clusters within a given operational service. The delivery was anticipated in Phases with a first system to be delivered in October 2008, replicated to make Phase 1 in January 2009 and upgraded in January 2011 to Phase 2. Some flexibility in delaying the start of service by up to 3 months and delaying the full implementation of Phase 1 by a further 3 to 6 months could be proposed. The evaluation would take the consequences of these delays into account.

Tender requirements covered mandatory, highly desirable and desirable features with a request for information. Performance requirements included time to solution of the benchmark codes involving multiple copies running concurrently. Running the benchmarks using hybrid programming results in the need for significantly less memory than pure MPI implementations. The requirements covered the need for pre and post-processing nodes specified as aggregate SPECfp rates, memory per node and total memory. Minimum usable

disk space and sustained bandwidth rates were also specified. Requirements for installation maintenance and support included minimum breaks during phase transitions, maintenance time per week, 24x7 hardware maintenance, on-site systems analyst, security requirements, system administration, development environment, scheduling, system monitoring, accounting and control, system availability and recoverability, physical environment, testing, migration, training and documentation, subcontracting, risk assessment, infrastructure costs and costs of delays. The ITT had some 330 Request for Information, some 320 Mandatory requirements, some 310 Desirable requirements and some 306 Highly Desirable requirements.

The timetable included:
- 13 April 2007: open briefing
- 1 June 2007: receipt of tenders
- June 2007: presentations by tenderers
- July-September 2007: demonstrations of tendered systems and evaluation of tenders
- October-November 2007: communication to tenderers of decision
- December 2007: contract signature

Acceptance Tests included functional tests to verify on-site that each Phase satisfies the technical specification and functional description (performance, memory, I/O and resilience), an operational test to demonstrate reliability and availability meets minimum requirements and a training/ test system. The service contract insists on on-site support

Benchmarks included the Integrated Forecast System – data assimilation 4D-var to meet commitments on memory and performance, the IFS forecast model and a set of kernels that measure I/O rates and network access rates. ASIS and optimised runs are requested.

ECMWF are currently running through their acceptance tests.

6.7 Procurement by NERSC, DoE

R.J. Blake – review of Request for Proposals

Introduction
The University of California through NERSC at the Lawrence Berkely National Laboratory is seeking to procure a balanced production scientific computing system that is a high performance parallel computational resource for the NERSC user community. The service has about ~ 3000 users and 400 projects – 10% of the time is allocated to INCITE high impact awards with a range of exploratory (10K hours) and broad impact (up to 5M hours) awards. A broad base of applications is supported including: particle physics, nuclear physics, chemistry, materials, engineering, plasma physics, geo-science, climate modelling and astrophysics. The performance goals include a Sustained System Performance (SSP) for the breadth of the scientific workload of 70-100TF average over the first 3 years. NERSC are looking for a balanced system (aggregate memory, global usable disk storage, bandwidth and latency, storage capacity and bandwidth, network bandwidth) and integration with the local environment including a global filesystem and external networks, adding best value to the local infrastructure.

Request for Proposals
The procurement process developed requirements through a Greenbook, analysis of usage patterns and through a series of workshops to provide a workload analysis and inform benchmarks/ test, minimum requirements and performance features. This was integrated with

a review of market constraints including technology, timing, budget and facility capabilities to inform the Request for Proposals.

The Request for Proposals was published on: September 19, 2008 with responses due by October 10, 2008. Delivery and installation is expected between August and October 2009 and acceptance of the system no later than December 2009.

NERSC actively embraces the PERCU method:

Performance – how fast will the system process work if everything is working really well.
Effectiveness – likelihood that users get the system to do their work when they need it.
Reliability – likelihood systems is available (to do users work)
Consistency – how often the system processes similar work correctly and in the same time.
Usability – how easy is it for users to process their work as fast as possible.

Minimum Requirements specifications included:

- Environment: Maximum power including cooling – phase/ voltage power requirements.
- Performance: min SSP, network (high bandwidth, low latency, fault tolerant, scalable interconnect), external network connectivity.
- Effectiveness: Application Development Environment, standards compliant F, C, C++, MPI(I.O), single application on all nodes.
- Reliability: comprehensive maintenance, 24x7 support for all HW and SW including replacement and spare parts.
- Consistency: correct, consistent and reproducible execution times in single user and production mode.
- Usability: IEEE 754 FP compliance, external file-system, administered from comp nodes, unified name space, highly parallel I/O, standards based and packaged independently, compatible with infrastructure.

Performance Features are assessed qualitatively in the various PERCU categories and include issues such as:

- Integrated system, multi-user, multi-application workload.
- Energy efficient computing, power distribution and cooling – standard thermal ranges/ changes.
- Ease and minimal cost of integration into existing infrastructure.
- Credible roadmap for future hardware and software and support thereof.
- Expandability, configurability of CPU, memory, storage and interconnect.
- Global shared storage system – 1 PB user space – 70 GB/s measurable sustained, aggregate bandwidth from file system to compute nodes.
- 100 GB/s network connectivity.
- Checkpoint-restart, job migration, backfill, reservations, pre-emption.

The procurement puts an emphasis on user viewpoints, uses PERCU to provide a holistic assessment for the system and seeks to deliver a system ready to run the entire production workload by the completion of acceptance. At the same time the procurement seeks to allow vendors to be innovative in technology, time and risk and in cost as a function of time.

Benchmarks

Benchhmarks include system component tests (Stream, PSNAP, multipong, IOR, Metabench, NetPerf), Kernels such as NPB, stripped down applications such as AMR Elliptic Solver, full

applications and composite tests SSP, ESP and CoV. Vendors are invited to submit base case results, optimised runs to highlight features of the system. The seven applications codes have different levels of concurrency up to 8096 processors, The benchmarks include code for emerging programming models and for new algorithms eg UPC, AMR and implicit sparse solves.

System performance is evaluated through the Sustained System Performance measure which is the geometric mean of the processing rates per core of seven applications multiplied by number of cores in the system. The Effective System Performance metric measures the efficiency in scheduling a defined workload through the system.

Evaluation of Proposals
NERSC embrace enthusiastically the Best Value Source Selection which allows vendors to propose their best solution given a set of minimum requirements and performance features. The Performance Features focuses on the strengths and weaknesses of the proposals within a framework of Feasibility governing likelihood of success, how well will the approach work, time risks, technological balance providing a timely, effective and manageable solution and Applicability (how well does the approach meet the procurement goals within a general user-environment and energy efficiency).

Suppliers are assessed on qualitative criteria in terms of their ability to produce and test the system, meet the schedule, provide continuous development, testing and support, corporate capability and risk management, commitment to HPC. These attributes are evaluated through a Capability framework in terms of advancing the state-of-the-art, past experience on similar projects, quality of project plans, management and personnel and interactions with third party software vendors.

Total life-cycle cost of ownership is increasingly important in terms of overall project Affordability.

Proposals will be evaluated by a panel of experts from LBNL. The successful proposal must meet all of the minimum requirements and contain the combination of performance features, supplier attributes and affordability offering the best overall value. The relative importance of all non-cost factors and life-cycle costs is approximately equal.

Contract Development
The RfP was issued to vendors and their responses will be used as a starting point for Statement of Works (SOW) negotiations. The vendor proposals will be evaluated, a single proposal selected and negotiations commenced. The SOW informs the development of the contract and performance metrics which governs system integration and acceptance testing. The Quality System for the procurement overviewed the steady state performance, maintenance arrangements and upgrades informed through performance monitoring. The goal is an excellent system for a long-time.

Acceptance
Acceptance tests require that the system is assembled and tested in situ including all hardware installation and assemble, burn in of all components, installation of software and approved production environment, tests and benchmarks addressing functionality, performance, reliability and quality and run benchmarks to demonstrate performance commitments.

In addition it is desirable to run a Factory Test prior to shipping covering power up and down, UNIX commands, monitoring software, reboot functions, power cycle from console, configuration testing benchmarks and an variability test expressed as some % over a number of days.

The acceptance test duration periods can be as long as 60 days requiring a 30 day contiguous user service.

Functionality demonstrations are run on the configuration that will go into production and include remote monitoring, power control and boot capabilities, network connectivity, filesystem functionality, batch system, system management software program development environment and UNIX functionality.

Systems tests include to production state in a reasonable time period and single node power tests. Performance tests are structured around benchmark suites ranging from simple component tests through to applications performance to throughput tests as discussed in the benchmarking section.

Availability tests will be run for 30 contiguous days in a sliding 60 day window requiring typically 98% availability –often running a selected user service. Failures include unavailability of nodes, inability to access the file system, inability to login, unavailability of full switch bandwidth, inability to launch batch operations,