eChallenges e-2012 Conference Proceedings
Paul Cunningham and Miriam Cunningham (Eds)
IIMC International Information Management Corporation, 2012
ISBN: 978-1-905824-35-9

The Strategy of the Commons: Modelling the Annual Cost of Successful ICT Services for European Research

Matti HEIKKURINEN [1], Sandra COHEN [2], Fotis KARAGIANNIS [2],
Kashif IQBAL [3], Eoin BRAZIL [3], Sergio ANDREOZZI [4]
[1]*Emergence Tech Ltd, St James House, 8 Overcliffe, Gravesend, Kent, DA11 0HJ, UK*
Tel: +41 (0)79 703 7347, Fax: +41 (0)86 0797037347, Email: matti@emergence-tech.com
[2] *Athens University of Economics and Business, Address, Athens , 104-34., Greece*
Tel: +302108203168, Fax: + 2108203164, Email:scohen@aueb.gr
[3] *Irish Centre for High End Computing (ICHEC),*
The Tower, Trinity Technology and Enterprise Campus, Dublin, D2, Ireland
Tel: +353 1 5241608, Fax: + 353 1 7645845, Email:{kashif.iqbal, eoin.brazil}@ichec.ie
[4] EGI.eu, *Science Park 140, 1098 XG Amsterdam, The Netherlands*
Tel: + 31 (0)20 89 32 007, Fax: + 31 (0)20 592 5155, Email: sergio.andreozzi@egi.eu

Abstract: The provision of ICT services for research is increasingly using Cloud services to complement the traditional federation of computing centres. Due to the complex funding structure and differences in the basic business model, comparing the cost-effectiveness of these options requires a new approach to cost assessment. This paper presents a cost assessment method addressing the limitations of the standard methods and some of the initial results of the study. This acts as an illustration of the kind of cost assessment issues high-utilisation rate ICT services should consider when choosing between different infrastructure options. The research is co-funded by the European Commission Seventh Framework Programme through the e-FISCAL project (contract number RI-283449).

1. Introduction

Assessing the actual costs of aggregate services consisting of multiple components from different sources is a difficult, but important task for several reasons. For example, the analysis forms the basis for assessing potential benefits of different outsourcing options and can help in managing pricing-related risks of Cloud-based services. The complexity stems from the need to take into account both capital and operational expenses (CAPEX and OPEX) of all of the components, which can vary based on local conditions and the actual level of use. Thus the apparent simplicity of the basic Cloud approach –only OPEX costs, "pay as you go" with a fixed unit price – often seems like an attractive, simple model. However, in reality Cloud pricing can include "CAPEX-like" components (such as billing for reserved instances) and other upfront costs cush as adapting applications to Cloud environments). Furthermore, as the volume of resource use grows, the premium pricing of the on-demand solutions may make them economically disadvantageous. We believe that in such a case any ICT service faces a similar challenge in finding the optimal mix of ICT infrastructure solutions as is currently faced by the European ICT services for research.

Dedicated High Performance Computing (HPC) and High Throughput Computing (HTC) services provided by initiatives such as PRACE[1] and EGI[2] support a wide range of

[1] Partnership for Advanced computing in Europe (http://www.prace-ri.eu/HPC-access).

basic and applied research activities in Europe. The original users – such as material sciences or High-Energy Physics researchers – already possessed fairly massive in-house ICT infrastructures. Thus the initial shift to service orientation came from the need to link computing centres together to solve problems that exceeded the capacity of an individual centre, which required comparatively modest "top-up funding" from the EC. What was essentially established is the "European e-Infrastructure commons", governed with flexible, often informal rules and conventions concerning the use of this common resource pool – a model that has been successfully applied for over a decade.

Little by little the use of this e-Infrastructure has extended to application areas where the ownership of high-end ICT resources has not been a prerequisite for the research. This has an impact on sustainability: estimating future resource use is harder and user new communities without their own ICT systems often require new services, as well as modifications to the practices that govern the use of the common resources[3]. The community needs also to adjust to new business models and technologies, such as ones brought on by Cloud computing. Until recently, most of the initiatives in this domain were fairly small pilot projects, but increased interest in Cloud interfaces, "Cloud federation" and use of commercial public Clouds as an e-Infrastructure component are bringing the issue to the forefront. The strategic approach needs to evolve to take new technical opportunities and cost models into account, while at the same time preserving the flexibility and ability to support large-scale, collaborative research projects in the long term.

2. Objectives

To support strategic planning, the e-FISCAL project has been funded to assess the overall cost of computing service provision for research in Europe. The project gathers together high-level experts in the e-Infrastructure service provision, cost assessment and policy development, also selected for their demonstrated ability to work in trans-disciplinary environments. The project uses a standard common metric, *cost per core hour*, as a way to assess basic cost efficiency of the different approaches. While this is an extremely reductionist, crude measure, it allows grounding sustainability issue in a way that is still relevant to funding agencies, researchers and ICT service providers. It is also considered to be one of the best tools to compare costs with cloud offerings [3].

However, despite its apparent simplicity, even "cost per core hour" is dependent on the application context. Different CPUs are optimised for different tasks, so there isn't a generic "cost per standard core hour". For this reason the cost analysis is complemented with a small-scale benchmarking exercise to gauge the possible impact this may have on the application-level performance. The project will also conduct interviews with different centres to better understand the rationale behind the choices made and different strategies (e.g. "bare" infrastructure vs. infrastructure *and* high-level user and developer support). Identifying typical service provision strategies for supporting different application domains might in turn point towards alternative metrics that could complement the core hour approach. In the long term, this analysis of the cost models and cost breakdown can be equally valuable results of the project. By supporting alignment of the cost categories, future pan-European cost assessment initiatives should be more accurate and require less effort. The cost models developed can be used in any context, but could be particularly useful in the public-private partnerships where the business and accounting models of the participants make the traditional cost modelling approaches particularly difficult to apply.

While the organisational structures of the HTC and HPC services are very well defined today, coming up with a consensus definition to separate them on the technology level is

[2] European Grid Initiative (http://www.egi.eu/)

[3] For example to accommodate parties that are not bringing in their own in-kind contributions

surprisingly difficult. At the moment we use the degree of parallelism as the differentiating factor: an HTC application may use the same inter-process communication mechanisms as an HPC one, but infrequently enough to allow use common, off-the-shelf networking solutions unlike a typical HPC application[4]. We also exclude GPU or FPGA based solutions from the scope of this cost comparison effort, as they were still niche approaches in the beginning of the data gathering period.

3. Methodology

The project applies a cost model – the e-FISCAL annual cost model – that avoids the shortcomings of the Total Cost of Ownership (TCO) and Full Cost Accounting analysis and permits the execution of high-level cost analysis. Complementary approaches, such as literature search and benchmarking are also used to complement survey data.

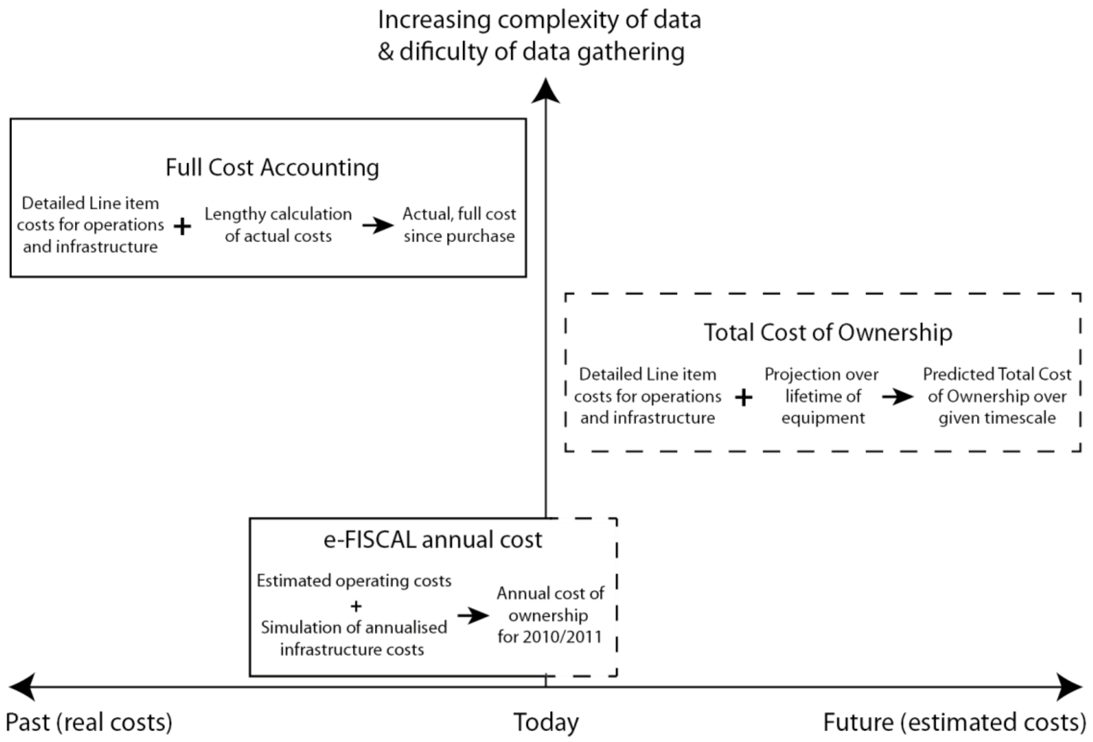

Figure 1: Comparing TCO, FCA and e-FISCAL methodology in cost calculations

TCO is a useful concept to assess the cost of a specific project over its useful life. Applying such a model adopts a *forward looking* stance, covering all *expected* costs over the project's lifetime that are divided by the anticipated useful life to come up with the yearly costs. However, to achieve acceptable precision of these predictions, several details need to be taken into consideration in a consistent manner - difficult task to accomplish across organisations. Nevertheless, this approach has been used extensively [1], [2], [3], and is often used as a planning and decision support tool.

Full Cost Accounting methodology relies on actual cost accounting data information available with the cost accounting systems of organizations, i.e. it adopts a *backward looking* stance. Detailed *actual* data (in line item format) is attributed and allocated based on various costing procedures to come up with the "cost per unit" of the object being analysed. Comparing this with the initial estimates (produced using e.g. TCO model) can be used to identify areas where more management attention is needed – TCO is mainly a tracking tool rather than a planning one. However, the reliability of the cost data depends

[4] While by no means the only difference, it provides a sufficient criterion for the study

on the robustness of the cost accounting system; and even with best accounting practices, costs that are not funded by the organisation are not registered as parts of the total costs.

In contrast, the e-FISCAL model faciliates accurate short to medium term estimates of the costs of maintaining services at their current level. It can do this without the need to identify funding sources or the exact times when the actual infrastructure investments have been made. Figure 1 presents the model compared to TCO and Full Cost accounting. Rather than relying on the detailed financial data or projections to the future, a relatively concise survey[5] about the infrastructure in place, recent unit acquisition and operational costs, makes it possible to estimate the overall annual cost of the infrastructure. This is more relevant and accurate for short-term planning purposes than Full Cost or TCO results. The survey can usually be answered based on information that persons in the operations management roles need to have at hand all the time (number of cores, basic description of the centre, recent budget numbers) without having to access, for example, archived financial data or make error-prone predictions decades into the future.

4. Technology Description

To put the model's cost information into the context of the actual research process, system benchmarking is used to compare and evaluate the performance. Even though limited in scope, this has an important secondary role in establishing a common overview model supporting trans-disciplinary approach of the project. Cloud services offer access to both HPC and HTC type systems and the benchmarks need to be selected accordingly. The benchmarks and test environments used in the project are presented below:

1. The NAS Parallel Benchmarks (NPB)[4] for HPC and HPC Cloud, designed to evaluate the performance of a typical HPC application. The suite consists of eight programs: five parallel kernels and three applications mimicking typical scientific software packages, all providing different problem sizes classes. For this study we have selected the Class B problem size[7] and use OpenMP[5] and MPI[6] versions.
2. The HEP-SPEC06 (based on the SPEC CPU2006), standard benchmark used by the HTC community. Due to space constraints HTC results will be published separately.

Table 1: HPC and HPC Cloud test environments

	Amazon EC2[6]	Stokes
Compute Node	23 GB of memory, 33.5 EC2 Compute Units (2 x Intel Xeon X5570, quad-core "Nehalem" architecture), 64-bit platform	24 GB memory, Two Intel (Westmere) Xeon E5650 hex-core processors per unit, 64-bit platform
Connectivity	10 Gigabit Ethernet	ConnectX Infiniband (DDR)
Compilers, Libraries	Intel C, Intel Fortran, Intel MKL, Intel MVAPICH2	Intel C, Intel Fortran, Intel MKL, Intel MVAPICH2

5. Developments

The e-FISCAL model workflow that forms the basis of the e-FISCAL cost assessment survey is presented in Figure 2. The identification of cost categories is principally based on [9] and formed the basis for making the unavoidable trade-offs in the questionnaire design (prioritising questions in order to reach sufficient response rate). The survey targeted National Grid Initiatives (NGIs, EGI country representatives), national HPC coordinators

[5] that still has enough detail and redundancy to provide accurate information

[6] We used the StarCluster [7] toolkit is used to create EC2 instances

[7] Middle-sized problem (presented in http://www.nas.nasa.gov/publications/npb_problem_sizes.html)

and also individual HTC/HPC centres. This was necessary, as some of the key information (such as energy or auxiliary costs) is usually available only at centre level. The survey data is used primarily for computing effective hourly cost per core hour and identification of major cost components. The development work has been reviewed based on extensive state-of-the-art study that also produced a publicly available reference list[16].

6. Results

6.1. Cost analysis

The analysis is based on the 26 answers from 14 countries,[8] belonging to NGI/EGI(12), the national HPC infrastructure or PRACE (2) – or both(7). The sample is rather evenly split between respondents providing computing services (CPU and storage) directly(13) and those providing both computing and coordination services(12). Table 2 presents the results in Euro per core for 2010 and 2011. The analysis uses average and median values based on valid answers received for each of the questions for all the parameters. As seen in Table 2, the initial cost estimates range from €0.033 (2011) to €0.10 per core hour (2010). However, the utilization rates reported in the calculations are rather low, and e.g. reaching a rate of 80% would bring down the cost to a range of €0.028 (2011) to €0.077 per hour (2010). This highlights the importance of the utilization rate (percentage of core hours actually used).

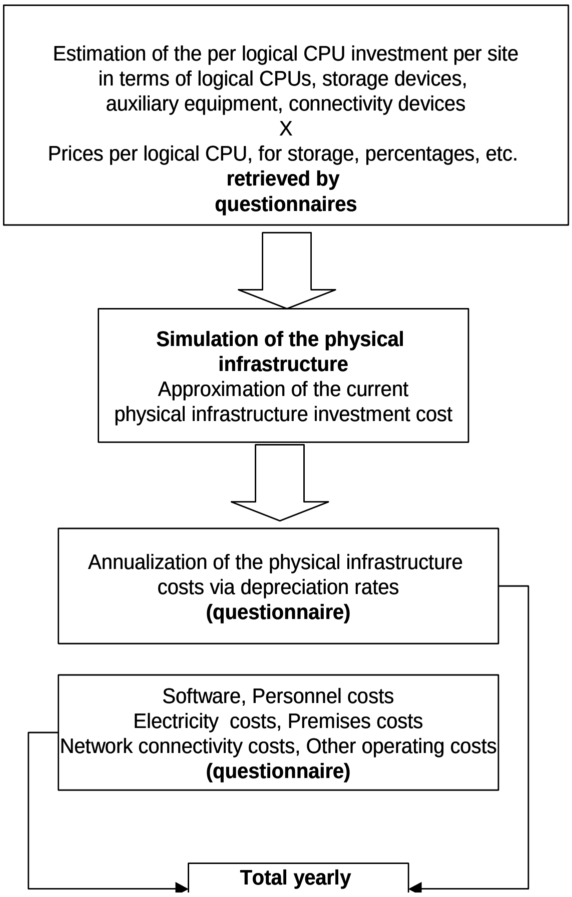

Figure 2: Overview of the e-FISCAL cost methodology to calculate total yearly cost

The acquisition costs of computing and storage devices decreased considerably from 2010 to 2011 (consistent with expectations), while investments that relate to interconnection devices and auxiliary equipment are more stable. The total yearly cost per logical core ranges from €197/year (2011, median) to €536 (2010, average).

The costs are close to the 0.05£ - 0.07£ per core hour discussed in [8], which is based on information shared in confidence by research computing centres in the UK. The numbers in the Table 2 are also in the same magnitude with the preceding cost assessment effort [9]. Some US reports present lower costs: the Hopper system operated by National Energy Research Scientific Computing Center reports $0.018/logical core hour[3] and Purdue campus grid $0.040[10], while with the Hyak cluster of the University of Washington [1] the cost estimate is 0.027$[9].

[8] Belgium, Bulgaria, Cyprus, Finland, Germany, Greece, Hungary, Ireland, Latvia, Norway, Poland, Romania, Spain and Turkey

[9] Based on reported yearly cost per node of $2,794, with a node having two six-core CPUs

Table 2: Preliminary e-FISCAL results (in €)

	2010		2011	
	Average	Median	Average	Median
Primary investment (CPUs + storage)/core	611	321	550	242
Other investment	148	67.	133	50.
Overall invested capital / core	759	388	684	292.
Total yearly **CAPEX**/ core	149	78	135	59
Total yearly operating costs **(OPEX)**/ core	387	180	288	139
Total yearly cost/ core	536	25	423	197
Utilization rate	59%	59%	58%	67%
Cost per core/hour	0.1036	0.0499	0.0837	0.0337

Apart from the aggregate information above, the data presented in Table 3 for 2010 and 2011 points towards areas for further study. Usually TCO method uses four (e.g. [11]) or three (e.g. [1] and [2]) years – as the useful lifetime of a CPU, the respondents used five years to calculate annual depreciation. This has an impact on the capital costs accounted each year in forming the total yearly cost. The average numbers of FTEs per core and the m^2 per core exhibit a decreasing trend that is consistent with economies of scale. This may be largely due to the new multi-core servers. Similarly, the electricity consumption per logical CPU shows a decreasing trend. According to [12] servers in 2010 (compared to 2005) have much higher processing power, more memory, faster network connections, more components and bigger power supplies. However, they also have improved power management that reduces electricity consumption. In cost models averages of €0.10/kwh and €100/m2 per year were used.

Table 3: Additional cost breakdown information derived from the analysis

	2010[10]			2011	
	Average	Median		Average	Median
FTEs/1000 core	6.54	2.80		4.71	2.08
Average cost per FTE in €	48,547	47,600		49,305	47,600
m2/'1000 core	80.93	59.58		60.38	49.34
Electricity consumption per core in kWh	379.41	363.41		284.80	307.74
Operating costs / total yearly costs	72.20%	69.89%		68.12%	70.27%
Capital costs / total yearly costs	27.80%	30.11%		31.88%	29.73%

The costs per core hour calculated above (including storage cost) could be compared with Amazon's E2C service prices[11]. However, detailed analysis is problematic, as both the costs and Cloud pricing are changing quite rapidly and comparing costs with prices is not straightforward (profit margin or loss-leader approach are hard to factor in). Nevertheless, in the short run costs should exhibit a greater stability. Nevertheless, continuous monitoring of both Cloud offerings and developments of the in-house costs is advisable before making strategic decisions. Scenario analysis should take into account that in research domain costs related to administration and application adaptation tasks will still remain even if the infrastructure itself would be fully outsourced.

[10]Investment relationships and depreciation rates were asked only for the most recent period (i.e. 2011).

[11]At the time of writing $0.09/hour (excluding storage and networking, on-demand small Linux instance)

6.2. Benchmarking results

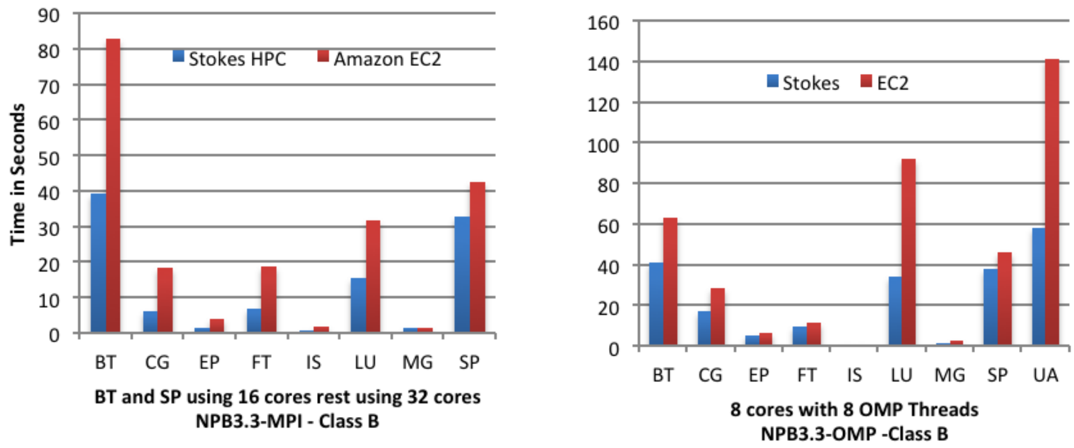

Figure 3: NPB MPI and OMP Benchmark Results

Figure 3 illustrates the benchmark findings for the MPI and OpenMP variants of the NPB. The MPI tests were run using three Amazon EC2 nodes and two Stokes nodes, whereas the OMP tests were run on a single node on both cases (eliminating the impact of the different networking solutions). The preliminary findings, which establish a basic range of potential performance normalisation factors, can be summarised as follows:

- For the OpenMP version of the benchmark, the purpose-built HPC system (i.e. Stokes) outperforms EC2 compute cluster for the same number of OMP threads. The average performance loss of moving from dedicated system to a Cloud was 37.26% (ranging from 16.18 - 58.93% for individual benchmarks).
- Also with the MPI versions of the NPB programs the EC2 performance lags behind the Stokes cluster. The likely differentiating element is EC2's 10 Gigabit Ethernet versus the InfiniBand interconnect used at Stokes. The average performance loss was 48.42% (ranging from 1.02% to 67.76%).

The above results show that even with a relatively modest problem size the performance differences between the virtualised general-purpose HPC infrastructure and a purpose-built system for HPC can be considerable. In addition to performance degradation, the configuration overhead for the EC2 is an additional factor that should be considered.

7. Business Benefits

The cost assessment can be used to assess potential financial implications of new technologies, since they will have a noticeable financial impact only if they influence major cost categories. The cost breakdown will also offer opportunities for optimisations on the management level: focusing the financial oversight to the areas where it will have most impact should allow developing a more streamlined organisational approach.

The value of the cost assessment is thus going to manifest itself especially as a part of the overall organisational development activities. Research activities and ICT solutions needed for them mature gradually to a level can be standardised and in some cases eventually supported by commodity solutions. Cost assessment can speed up the process and release more resources to the development of new e-Infrastructure innovations. Outside the research domain rapid and low-cost cost analysis helps identifying areas where considering alternatives to the basic commodity technologies may be cost effective. The models developed may help development of contractual arrangements to manage the risks related to the "pricing vs. actual costs" issue that is inherent to any outsourced solution.

On more abstract level, the critical role of e-Infrastructure in many of the EC strategic flagship initiatives means that the sustainability of the services needs to be ensured – also in situations where individual countries or regions face serious economic difficulties. Designing Pan-European approach for the sustainability that can answer to these challenges is possible only if it is based on accurate estimates of the costs of the service provision.

8. Conclusions

It should be noted that any survey – especially if they are among the first of their kind – often requires considerable outreach and engagement efforts before producing results. Even Top-500, well established survey today, faced this issue[12]. The sensitive, strategic nature of financial information accentuates this, necessitating consistent communication and trust building. In addition to confidentiality of the data, one of the key issues is being open about the limitations of the approach: Top-500 benchmark results or reported cost per core hour do not capture information about the value or impact of the activities measured.

Being able to accurately break down the costs is a key prerequisite for going beyond basic budgetary sustainability planning and optimisation. As the research use of ICT becomes more multifaceted, being able to identify and assess the costs of the minimum resource requirements (for which a dedicated infrastructure is well-suited) as well looking into new technical and financial options to support the burst-like peak activity more efficiently (often a good match with outsourced solutions) is important.

From the socioeconomic point of view knowing the costs of different ICT service provision models is only the first (albeit crucial) step. Understanding how the activities supported by ICT solutions actually produce value, and studying the different tangible and intangible effects the choice of the provision model has on this process are important topics for further research. Making drastic changes in the research HTC and HPC service provision and their funding principles without this knowledge would – even in the best case – be imprudent and most likely have results that are counter-productive.

References

[1] University of Washington "Hyak Operating Costs and Comparison with Commercial Alternatives", http://escience.washington.edu/get-help-now/hyak-operating-costs-and-comparison-commercial-alternatives (accessed 18/07/2011)

[2] Nazir, A, S- A. Sørensen (2010), *"Cost benefit analysis of high performance computing infrastructures"*, Presented in Service Oriented Computing and Applications (SOCA) 2010 IEEE International Conference.

[3] Magellan Final Report, December 2011 30. (http://science.energy.gov/~/media/ascr/pdf/program-documents/docs/Magellan_Final_Report.pdf)

[4] NAS Parallel Benchmarks: http://www.nas.nasa.gov/Resources/Software/ npb.html (accessed 25/05/2011)

[5] OpenMP specification: http://openmp.org (accessed 25/05/2012)

[6] Message-Passing Interface (MPI) specification: http://www.mpi-forum.org/.

[7] http://web.mit.edu/star/cluster/ (accessed 25/05/2012)

[8] Hawtin, R. M. Hammond, L. Gillam and G. Curtis, "Cost analysis of cloud computing for research", Final Report to EPSRC and JISC, 22/2/2012, Curtis+Cartwright Consulting Ltd.

[9] e- IRGSP2 EU deliverable D4.3 - Final legal issues report public, e-IRGSP2 Consortium, 28/01/2011

[10] Smith, Preston M. M.S., Purdue University, May 2011. A Cost-Benefit Analysis of a Campus Computing Grid . Major Professors: Jeffrey J. Evans, Kevin C. Dittman, and Carol X. Song.

[11] Walker, E. (2009), "The real cost of a CPU hour", *Computing Practices*, Volume 43, Issue 4, pp. 35-41.

[12] Jonathan Koomey. 2011. *Growth in Data center electricity use 2005 to 2010*. Oakland, CA: Analytics Press. August 1.

[13] The e-FISCAL project state-of-the-art repository: http://www.efiscal.eu/state-of-the-art (accessed 26/05/2012)

[12] See http://crd.lbl.gov/news-and-publications/news/2012/erich-strohmaier-and-the-top500-a-list-that-s-taken-on-a-life-of-its-own/

