



Financial Study for Sustainable Computing e-Infrastructures

Deliverable Title: D2.2 Computing e-Infrastructure cost calculation at national and European level

Partner Responsible: AUEB-RC

Work Package 2: Computing e-Infrastructure cost calculation

Submission Due Date: 31.05.2012

Actual Submission Date: XXXXX

Distribution: Public

Nature: Report



Abstract: This report summarizes the preliminary findings from computing e-infrastructure costs calculations. It includes a description of the collection and synthesis of questionnaire data as well as the corresponding preliminary analysis in order to come up with quantitative and qualitative results such as the cost per logical CPU/hour. The deliverable also includes a brief overview of the current state-of-the-art and a section about benchmarking.

Document Information Summary

Deliverable number:	<i>D2.2</i>
Deliverable title:	<i>Computing e-Infrastructure cost calculation at national and European level</i>
Editor	<i>Sandra Cohen</i>
Contributing Authors:	<i>Sandra Cohen, Fotis Karagiannis, Costas Courcoubetis AUEB-RC, Matti Heikkurinen ETL, Kashif Iqbal ICHEC</i>
Reviewer(s):	<i>Sergio Andreozzi, EGI</i>
Work package no.:	<i>WP2</i>
Work package title:	<i>Computing e-Infrastructure cost calculation</i>
Work package leader:	<i>Sandra Cohen</i>
Work package participants:	<i>Sergio Andreozzi, Owen Appleton, Eoin Brazil Costas Courcoubetis, Matti Heikkurinen, Kashif Iqbal, Fotis Karagiannis, Steven Newhouse</i>
Distribution:	<i>Public</i>
Nature:	<i>Report</i>
Version/Revision:	<i>v.final</i>
Draft/Final	<i>Final</i>
Keywords:	<i>e-Infrastructure costs, Total yearly cost of ownership, TCO, questionnaire, Cost per logical CPU/hour, High Throughput computing, High Performance Computing, Cloud computing, HTC and HPC cost break down, Benchmarking, State-of-the-art, Hyper-threading, PUE, Green IT, Personnel, CAPEX, OPEX</i>

DISCLAIMER

e-FISCAL is a Research Infrastructure Project co-funded by the GÉANT & e-Infrastructure Unit of the Information Society & Media Directorate General of the European Commission. e-FISCAL targets computing e-Infrastructure providers, national funding agencies, scientific communities, as well as European Union policy makers. e-FISCAL is supported by an external Advisory Board (AB). This deliverable has not been reviewed by the external Advisory Board.

This document has been produced with the co-funding of the European Commission. The content of this publication is the sole responsibility of the e-FISCAL consortium and cannot be considered to reflect the views of the European Commission.

Grant Agreement Number: 283449

Start Date: 1 August 2011

Duration: 18 Months

Log Table

Issue	Date	Description	Author/Partner
v0.1	17/4/2012	Deliverable outline- Skeleton	Sandra Cohen - AUEB
V0.2	25/4/2012	Deliverable outline- Final Skeleton	Sandra Cohen and Fotis Karagiannis - AUEB
V1.0	31/5/2012	Deliverable 2.2. – First draft	Sandra Cohen and Fotis Karagiannis - AUEB
v2.0	8/6/2012	Deliverable 2.2. – Second draft	Sandra Cohen, Fotis Karagiannis AUEB, Kashif Iqbal - ICHEC
v3.0	15/6/2012	Pre-final version for internal review	Sandra Cohen, Fotis Karagiannis, George Thanos, Costas Courcoubetis AUEB
v3.1	19/6/2012	Internal reviews and additions	Sergio Andreozzi - EGI.eu, Matti Heikkurinen – ETL, Fotis Karagiannis - AUEB
v3.2	25/6/2012	Integration of comments from PMB members	Sandra Cohen, Fotis Karagiannis - AUEB
v3.3	26/6/2012	Further comments and contributions from PMB and their integration	Matti Heikkurinen – ETL, Sandra Cohen, Fotis Karagiannis – AUEB
v.final	28/6/2012	Final version	Sandra Cohen, Fotis Karagiannis - AUEB

Document Approval

Approval by	Date
Project Management Board	XXXXX

EXECUTIVE SUMMARY

Deliverable 2.2 (Computing e-Infrastructure cost calculation at national and European level) summarizes the preliminary findings of the e-FISCAL project, focusing on the cost estimations of the European computing e-infrastructure based on questionnaire and other relevant data. It includes a description of the collection and synthesis of questionnaire data, as well as the corresponding preliminary analysis in order to come up with quantitative results regarding the cost per logical CPU/hour. The results are also compared with other findings from Europe and US that are presented as part of the state-of-the-art review. First indications on qualitative analysis (studying parameters such as utilization) are also included. The preliminary efforts of the project on benchmarking in-house and commercial infrastructures are briefly summarized; at this first step the project compared an in-house developed and operated HPC infrastructure (ICHEC HPC centre in Ireland) with a leased HPC service in Amazon (being the dominant player in the IaaS market). Several best practices from the area of Green IT are also included in the deliverable. It should be noted that two versions of this deliverable have been prepared, a public version with anonymous summaries at European level and a restricted version with detailed information on the respondents and the corresponding centres included in additional annexes.

The main findings of the project have been the following:

- State-of-the-art review: There is a lot of recent literature on the costs of individual “in-house” developed HPC or HTC centres or campus Grid systems and their comparison with commercial “on-demand” cloud services, and in particular the dominant player in the market, i.e. Amazon Elastic Compute Cloud (Amazon EC2). The majority of these studies use a price for core/hour which ranges from around 0.015 Euros (Magellan report) to around 0.40 Euros (University of Washington) in the US¹; around 0.075 in the UK (Hawtin et al. 2012 for JISC) and around 0.09 Euros for a sample of EGI centres in Europe (e-IRGSP2 small scale study in 2009).
- The sample for the e-FISCAL findings was relatively good; 26 answers from 14 countries. However, high-end HPC centres (such as the PRACE Tier-0s) or other high-end HTC centres (such as the WLCG Tier-1s participating in EGI) are not included. This was mostly due to confidentiality reasons, particularly specific non-disclosure agreements between the vendors and those centres that do not allow for the publication of detailed cost information. There are on-going discussions with the coordinators of the EGI-InSPIRE and PRACE IP projects aiming to gather a minimum degree of information from these centres by the end of the project (through a “lighter” questionnaire). Striking a balance between the level of detail and ease of use was challenging, but we had to go for a relatively “ambitious” one in the first year. PRACE is also working on the estimation of Total Cost of Ownership (TCO) of its infrastructure, and we hope to be able to use the calculations as additional input.
- The e-FISCAL median values (which take into account the so-called “outliers” that are either very high or low compared to typical values) are around 0.05€/CPU hour in 2010 and 0.0337€/CPU hour

¹ All the exchange rates are recent i.e. June 2012

in 2011, while averages are similar to the e-IRGSP2 ones (around 0.10 €/CPU hour in 2010 and 0.08 in 2011). This shows that the e-FISCAL initial findings are in-line with studies reported elsewhere.

- The significant differences between median and average show that there are outliers in the sample (low or high numbers) that significantly influence the averages. These will be studied and the survey respondents will be asked during the e-FISCAL workshop or off-line to justify their values.
 - Costs between 2010 and 2011 are decreasing in-line with the trend of lower hardware prices and better overall efficiency.
 - The breakdown between CAPEX and OPEX in 2011 in our calculations is around 30%-70% (median) to 32%-68% (average). Around of 50-55% of total costs is dedicated to personnel costs. It is obvious that the personnel costs for a very large computing centre (in the order of 100.000 cores) can show economy of scale if compared to the same capacity distributed in smaller sites that are federated together. Further analysis will explore this issue in more detail.
 - The average utilization rate used to calculate the average and median cost per logical/CPU hour for the above results for 2011 is 58% and 67% respectively. This refers to a mixture of EGI, PRACE and other sites not integrated in these e-Infrastructures. As an example, for 2011, EGI reports an utilisation rate of 71,3%. The utilization rate in the e-FISCAL project has been calculated by taking into account yearly logical wall clock time and available logical CPUs at the end of the period (either 31/12/2010 or 31/12/2011). This assumes that the number of logical CPUs reported at the end of the period is available throughout the year. If this is not the case (e.g. because of an infrastructure upgrade towards the end of the year), the utilization rate calculated is underestimated. Obviously the higher the utilization, the lower the cost, but in the initial analysis “conservative” values were selected consistently.
 - Other interesting findings are the high numbers of depreciation rates for the hardware (average 5 years), the quite good rates of PUE efficiency (of around 1.5 median value) and the percentage of electricity cost (around 15% median value of all costs).
- Comparing with commercial on-demand prices such as Amazon EC2 who is the market leader is not straightforward due to several reasons: there is no performance normalization (benchmarking “sanity check” efforts are on-going, but not yet concluded), network and storage costs models for Amazon are different from dedicated infrastructure’s (N.B. in general the network cost information was not the primary focus of the e-FISCAL questionnaire), the fact that personnel (such as application developers and administrators) for configuring and operating EC2 instances and for the adaptation of the application code would be needed if researchers migrate to Amazon, and the inherent uncertainty of comparing in-house costs with EC2 prices² (not to mention costs from 2010 and 2011 with prices in 2012) are some of them.
 - Benchmarking HPC, HTC and commercial Cloud (Amazon EC2) performance is a small-scale effort in the project and is acting as a “sanity” check. As a first step, the NAS Parallel benchmark has been run

² Pricing is influenced by the profit margins, which in some cases can also be negative.

in both an HPC system (i.e. Stokes HPC cluster in Ireland) and a compute cluster instance from EC2. The results have highlighted that even for relatively moderate problem sizes the performance degradation is significant, averaging around 40%. Therefore, additional cost factors in terms of performance penalties and configuration overhead should be considered while estimating the cost for various e-infrastructures. The next step will use a similar test case comparing HTC infrastructures and corresponding Cloud instances. This will identify additional cost factors for the HTC and Cloud infrastructures.

It should be noted that during the course of the project, the issue of comparison between a physical core and a logical core (the OS abstraction of a physical core) was identified as warranting further study. In case of hyper-threading the operating system assigns two logical cores (virtual cores) for each physical core and shares the workload between them when possible³, however the terminology and resource monitoring practices are not fully unambiguous. As the questionnaire included the term “logical CPUs” and referred to the “number of processing cores”⁴, checking whether the respondents of the questionnaire have provided their answers based on physical or logical cores (i.e. taking into account or not hyper-threading) is important. The e-FISCAL workshop in Samos will provide an excellent opportunity for these checks, as well as allowing understanding and categorising the different monitoring practices in a way that minimises the chances of confusion.

³ Hyper-threading: <http://en.wikipedia.org/wiki/Hyper-threading>

⁴ The intention was to check for the number of physical cores

EXECUTIVE SUMMARY	4
1. INTRODUCTION.....	10
1.1 Scope of the Document.....	10
1.2 Target Audiences.....	10
1.3 Structure of the document.....	10
1.4 Terms and definitions.....	11
2. COST INFORMATION COLLECTION PROCESS	13
2.1 Review of the methodology used in the project	13
2.2 Questionnaire dissemination.....	15
2.3 Questionnaire follow-up	16
2.4 Summary of data collected	16
3. LITERATURE REVIEW	22
3.1 Overview of the state of the art.....	22
3.2 Conclusions retrieved by the state of the art review.....	25
4. DATA ANALYSIS.....	27
4.1 Cost analysis – Presenting the input values	27
Logical CPU information.....	27
Storage information	27
Cost information	28
Depreciation rates	29
Cost relations over investment.....	29
Auxiliary equipment information.....	30
Personnel information	31
Space and PUE information	31
Cost analysis – Metrics calculation	32
4.2 Calculation of the logical CPU cost/hour.....	35
4.3 Sensitivity analysis.....	37
5. BENCHMARKING	41
5.1 Background and Motivation	41
5.2 System Benchmarks	41

The NAS Parallel Benchmark (NPB).....42
The HEP-SPEC benchmark.....42

5.3 System Specification43
Stokes HPC System.....43
Amazon Elastic Compute Cloud (EC2).....43

5.4 NPB Execution and Results45

5.5 Summary and Future Work46

6. CONCLUSIONS AND FUTURE PLANS..... 47

6.1 CONCLUSIONS.....47
6.2 FUTURE PLANS.....49

REFERENCES & RESOURCES..... 50

TABLES

Table 1 – Countries contributing to e-FISCAL survey.....	17
Table 2 – Questionnaire overview- cost categories and other related information ...	19
Table 3 – Types of services represented in e-FISCAL survey.....	19
Table 4 – Institute’s infrastructure participation represented in e-FISCAL survey	19
Table 5 – Institute’s roles represented in e-FISCAL survey.....	21
Table 6 – Literature review summary in relation to Cost per logical CPU/hour or Cost per core/hour.....	26
Table 7 – Logical CPU information.....	27
Table 8 – Storage information	28
Table 9 – Acquisition cost of logical CPU and storage information.....	28
Table 10 – Useful life information.....	29
Table 11 – Cost relations over investment	30
Table 12 – Auxiliary equipment information	31
Table 13 – Salary information.....	31
Table 14 – Site information.....	32
Table 15 – Yearly Cost per logical CPU (in €)	34
Table 16 – Operating Yearly Cost per logical CPU (in €).....	35
Table 17 – Cost per logical CPU/hour (in €).....	35
Table 18 – Cost per logical CPU per cost category (in €) and % for 2010.....	36
Table 19 – Cost per logical CPU per cost category (in €) and % for 2011.....	37
Table 20 – Cost per logical CPU / hour (in €) under alternative scenarios	37
Table 21 – Allowed use of founding for buying cloud computing services	38
Table 22 – Using and buying cloud computing services in 2011.....	38
Table 23 – Using and buying cloud computing services from 2012 onwards	39
Table 24 – Green IT related actions in 2011 and future prospects.....	39
Table 25 – Sustainability related information	40
Table 26 – Class B Program Characteristics.....	42
Table 27 – Stokes and EC2 - System specification.....	44

FIGURES

Figure 1: Overview of the proposed methodology.....	13
Figure 2: Comparing TCO, FCA and e-FISCAL methodology in cost calculations	14
Figure 3: Map of participating countries	17
Figure 4: Graphical representation of e-FISCAL methodology.....	33
Figure 5: NPB MPI and OMP Benchmark Results.....	45

1. INTRODUCTION

1.1 Scope of the Document

The scope of this document is to present the preliminary findings from computing e-infrastructure national costs calculations. It includes a description of the collection and synthesis of the questionnaire data. Descriptions of the characteristics of the HPC and HTC centres that have contributed to the study through completing the e-fiscal questionnaire are presented. This report details the preliminary analysis of the input data. It acts as a useful means for information validation; as these preliminary findings are intended to be used as the basic material to be discussed in-depth with contributors in order to gather feedback. The goal of this feedback and the validation procedure is to reveal the weak and strong points within the costing methodology and related analysis process, and to conclude on viable ways to improve the accuracy and comprehensiveness of the results. Moreover, in this report a preliminary computation of the average cost per logical CPU/hour for the sample sites is performed. Albeit the extensive analysis and the thorough presentation of input data the specific information presented by each respondents is in no case disclosed in order to preserve confidentiality. This deliverable provides a brief overview of the current state-of-the-art and a section about benchmarking. Finally, the report describes a number of Green IT initiatives of the sample institutes and their stance toward cloud commercial offerings.

1.2 Target Audiences

The document is intended as an internal and external reference: internally as the main tool for the evaluation against the objectives in the DoW and the related metrics; external, as the core of this document is public and may be of interest of a series of stakeholders, namely:

- e-Infrastructure computing providers and HPC/HTC centre managers, providing them with a better understanding of the cost structures and actual costs of in-house vs. leased infrastructures to assist in planning and determining their centres' sustainability
- Research communities that can benefit from cloud computing allowing them to compare their costs against in-house infrastructures
- e-Infrastructure policy makers and funding agencies, who recognise the cost parameter as key parameter in the sustainability of computing e-Infrastructures

1.3 Structure of the document

The structure of the document is as follows: In Section 2 the cost information collection process is discussed. More specifically, an overview of the survey methodology followed so far in order to assess the cost of HTC and HPC centres for 2010 and 2011 is presented emphasizing on the survey dissemination and follow up procedures. In Section 3 a brief review of the state of the art is presented as well as the main conclusions that can be derived from its synthesis. Section 4 presents an analysis of the input retrieved from the questionnaires. Each questionnaire question is discussed and some summary metrics are presented. Comparisons with earlier empirical works, data from the EGI compendium, as well as comments raised by respondents are discussed throughout the relevant chapters. The report discusses the gathered data per identified cost category in detail, but presents only aggregate information for confidentiality reasons. Shortcomings of data as well as issues that deserve a more in-depth investigation are presented. Approximations of the yearly cost of a logical CPU as well as

the cost of logical CPU/hour are attempted on the basis of hypotheses backed on average and median resource values derived from the survey sample. The last part of Section 4 is devoted to data gathered from respondents regarding cloud computing, Green IT initiatives and sustainability. Section 5 is devoted to benchmarking. The report conclusions are presented in Section 6.

1.4 Terms and definitions

Cloud computing	Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.
EGI – European Grid Infrastructure	A federation of shared computing, storage and data resources from national and intergovernmental resource providers that delivers sustainable, integrated and secure distributed computing services to European researchers and their international partners.
EGI.eu	A non-profit organisation based in Amsterdam established to coordinate and manage the infrastructure (EGI) on behalf of its participants: National Grid Initiatives (NGIs) and European Intergovernmental Research Organisations (EIROs) .
HPC- High Performance Computing	HPC is a computing paradigm that focuses on the efficient execution of compute intensive, tightly-coupled tasks. Given the high parallel communication requirements, the tasks are typically executed on low latency interconnects which makes it possible to share data very rapidly between a large numbers of processors working on the same problem. HPC systems are delivered through low latency clusters and supercomputers and are typically optimised to maximise the number of operations per seconds. The typical metrics are FLOPS, tasks/s, I/O rates.
HTC - High Throughput Computing	HTC is a computing paradigm that focuses on the efficient execution of a large number of loosely-coupled tasks. Given the minimal parallel communication requirements, the tasks can be executed on clusters or physically distributed resources using grid technologies. HTC systems are typically optimised to maximise the throughput over a long period of time and a typical metric is jobs per month or year.
Hyper-threading⁵	Hyper-threading (officially Hyper-Threading Technology or HT Technology, abbreviated HTT or HT) is Intel's term for its simultaneous multithreading implementation first appearing in February 2002 on its Xeon server processors and in November 2002 on its Pentium 4 desktop CPUs. Later, Intel included this technology in Itanium, Atom, and Core 'i' Series CPUs, among others. Intel's proprietary HT Technology is used to improve parallelization of computations (doing multiple tasks at once) performed on PC microprocessors. For each processor core that is physically present (physical core), the operating

⁵ <http://en.wikipedia.org/wiki/Hyper-threading>

	system addresses two virtual or logical cores, and shares the workload between them when possible. The main function of hyper-threading is to decrease the number of dependent instructions on the pipeline.
IaaS- Infrastructure as a Service	One of cloud computing models; in this most basic cloud service model, cloud providers offer physical or more often as virtual computing cores (virtual machines), storage, network access and other infrastructure resources. IaaS providers supply these resources on demand from their large pools installed in data centres.
Logical processor or logical CPU⁶ or logical core	A processor (CPU) that handles one thread of execution (instruction stream). A logical processor can be a (physical) core or a hyper-thread. There can be one or more logical processors per (physical) core (more than one if hyper-threading is enabled) and one or more cores per processor socket.
NGI - National Grid Initiatives	NGIs are the entities responsible of procuring and operating the national grid infrastructure (in terms of computers and storage devices) and corresponding services to the research and academic communities.
PRACE (Partnership for Advance Computing in Europe)	PRACE is a unique persistent pan-European Research Infrastructure for HPC implementing 3-5 petaflop supercomputing systems in Europe. PRACE manages extreme computing power and a selected set of highly specialized services.

⁶ <http://msdn.microsoft.com/en-us/library/dd722831%28v=bts.10%29.aspx>

2. COST INFORMATION COLLECTION PROCESS

2.1 Review of the methodology used in the project

The methodology defined in this project in order to calculate the cost of the HTC and HPC centres for 2010 and 2011 comprises of six stages:

- State of the art review in costing issues
- Development of the generic cost model
- Sample selection
- Questionnaire development
- Questionnaire dissemination and follow up
- Analysis of data and conclusions

A graphical representation of the sequence of these steps is shown in Figure 1.

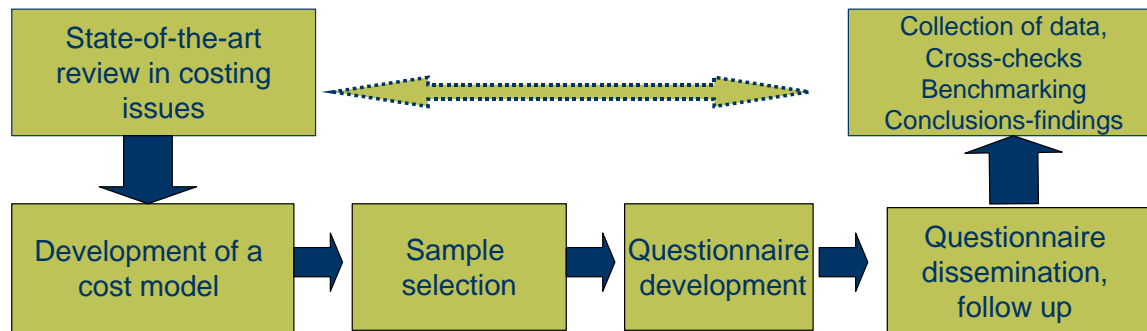


Figure 1: Overview of the proposed methodology

A presentation of the first four steps has been given in D2.1 e-FISCAL questionnaire. However, for purpose of completeness we briefly discuss these steps also here.

After reviewing the relevant state-of-the-art and by bearing in mind the purpose of the e-FISCAL project, we concluded that it could be better to develop a hybrid methodology that builds on the Total Cost of Ownership (TCO) and Full Cost Accounting while avoiding their shortcomings (see section 3) to meet our needs. The methodology developed allows getting accurate short- to medium-term estimates of the costs needed to maintain the HPC and HTC services at their current level and permits the execution of high-level cost analysis. It also allows calculating these costs without the need to identify funding sources or the exact times when the actual infrastructure investments have been made. A graphical representation of the e-FISCAL model compared to TCO and Full Cost accounting follows (Figure 2).

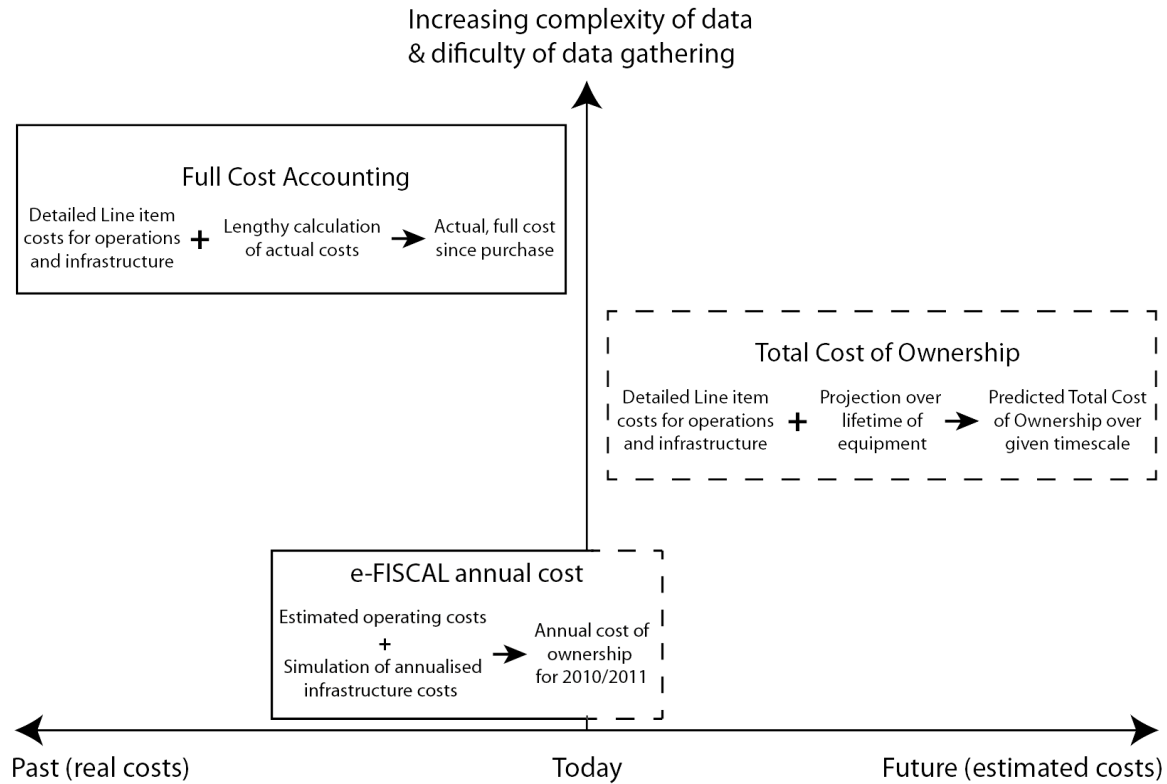


Figure 2: Comparing TCO, FCA and e-FISCAL methodology in cost calculations

The e-FISCAL costing methodology is materialised in a two-step process: a) Simulation of the physical infrastructure and b) Development of the financial model. The (a) corresponds to the investment cost of the infrastructure which is approximated by taking into account the capacity in terms of logical CPUs, of storage devices, of interconnection devices and of auxiliary equipment approximated on the basis of actual purchase values corresponding to each specific site/centre. The (b) is the sum of the annualized cost of the simulated physical infrastructure (Capital Expenses -CAPEX) and the operating cost of the physical infrastructure (Operating Expenses- OPEX). The annualized cost of the simulated physical infrastructure is derived by applying depreciation rates to annualize the cost of the physical infrastructure simulated in the first step. The operating cost of the physical infrastructure corresponds to the yearly costs for running the site/centre (e.g. personnel, electricity, premises cost).

The questionnaire developed for the purpose of the study contains proper questions to feed in the model with the necessary quantitative information. The distinct cost categories used in the fourteen sections of the questionnaire are found in literature (see for example Opitz et al., 2008).

In the following paragraphs information about the questionnaire dissemination and follow up is discussed. Section 4 is devoted to collection of data and preliminary findings presentation.

2.2 Questionnaire dissemination

The project made a press release to announce the initiation of the survey⁷. The press release was launched through several channels (such as Cordis Wire and AlphaGalileo), and was picked up by publications HPCWire, iSGTW, RoadRunner⁸, DallasNews⁹ and OSSREA¹⁰, among others. At the same time the questionnaire was promoted by e-mails to personal contacts. This proved quite efficient, as potential respondents could identify their past or current collaborators behind the questionnaire effort. The survey instrument was available in both on-line format (SurveyMonkey software¹¹) as well as a pdf editable format. As stated on the invitation, the questionnaire was soliciting answer from NGIs, national HPC coordinators, but also from individual HTC/HPC centres. This was deemed necessary as some costs or related information is kept only at this level (e.g. energy consumption or housing costs, especially if a national entity is not paying for these).

The initial deadline for questionnaire completion was the 14th of February 2012. However, an extension was given for the end of February due to several requests. The extension was officially announced on the project web site (<http://www.efiscal.eu/survey-extension>). In parallel with the extension, a letter from the European Commission GÉANT and e-Infrastructures Unit Head (Kostas Glinos) was sent to the EGI and PRACE Councils^{12,13} along with the e-IRG delegates¹⁴, putting the relevant weight and also explaining that dedicated effort for filling in a demanding questionnaire would have been non-economic. The letter was catalysing besides specific upscale HPC centres (such as the PRACE Tier-0 centres) that had binding non-disclosure agreements with the hardware vendors and could not reveal their costs. For these cases discussions with the PRACE IP projects coordinator (Thomas Eickermann) have taken place, who stated that a TCO analysis for these centres is taking place in the first semester of 2012 and some preliminary results will be available during summer. In addition, the nominated PRACE member in the Advisory Board, Pascal Moussier, and after June 2012 Annaig LeGuen, has been also trying to get some high-level data from France (GENCI). It is currently unclear what level of data the project is going to get for the upscale HPC centres, however on the second year of the project there will

⁷ The announcement is also found on the e-fiscal site (<http://www.efiscal.eu/node/15>).

⁸ <http://www.rr.com/>

⁹ <http://topics.dallasnews.com/>

¹⁰ <http://www.ossrea.net>

¹¹ SurveyMonkey: www.surveymonkey.com

¹² EGI Council members: http://www.egi.eu/about/EGI.eu/council_members.html

¹³ PRACE members: <http://www.prace-ri.eu/Members>

¹⁴ e-Infrastructure Reflection Group members: <http://www.e-irg.eu/about-e-irg/members.html>

be effort in this direction, along with a foreseen workshops co-organised with the EGI and PRACE projects.

2.3 Questionnaire follow-up

For questions related to the questionnaire a dedicated e-mail addressed has been set up (survey@efiscal.eu). Additionally respondents can send their queries though the contact page of the project.

Through informal discussions as well as through e-mail exchange it was evident that several people had concerns about the confidentiality of their responses as well as the potential use of the information provided for purposes outside the scope of the study. In order to dissolve any doubts about the confidentiality of input data, we developed a sample presentation with dummy figures in order to show how we indented the results to be reported¹⁵. In this presentation we made clear that the data provided will be processed in such a way that it will not be possible to identify the provider of any type of information. We also developed a more formal document outlining the confidentiality practices of the project that could be used as a basis for a signed agreement between the project and the party providing the information. However, at the time of writing the issues related to confidentiality have been solved without having to resort to signed agreements, sending the document outlining the confidentiality issues has sufficed in cases where the issues have not been insurmountable.

Finally, in order to stimulate participation and to ensure that the next project workshop (aiming to present this report, meet with the respondents and receive their feedback) would gather maximum number of financial experts, we pledged that those being the first to fully complete the questionnaire would be invited to the e-FISCAL financial experts' workshop in summer 2012. Additionally, they would be entitled to have their travel expenses to attend the event covered.

2.4 Summary of data collected

The majority of answers came through the SurveyMonkey software. Out of the 19 countries that had shown interest in participating in the e-FISCAL project (through visiting and browsing through the on-line questionnaire) we have representatives of 14 countries in our final sample.

Therefore this initial set of responses and interest in the survey almost satisfies one of the target values set at the outset of the project that of having more than 50% of the countries (NGIs or National HPC centres) compared to the invited NGIs ones participating in the study (19 counties out of the 33 countries participating to EGI.eu have shown interest while 14 out of the 33 completed the questionnaire)¹⁶. However, we hope to get another at least 3 countries by the end of the project.

The breakdown of the contributions per country is presented in the following figure and table.

¹⁵ This presentation is found in <http://www.efiscal.eu/files/sample%20publication%20eFISCAL.pdf>.

¹⁶ The National Grid Initiatives (NGIs) are organisations set up by individual countries to manage the computing resources they provide to the European Grid Infrastructure (EGI). Information about the NGIs is found in <http://www.egi.eu/about/ngis/index.html>.



Figure 3: Map of participating countries

	Country Name	Number of questionnaires
1	Belgium	5
2	Bulgaria	1
3	Cyprus	1
4	Finland	1
5	Germany	1
6	Greece	4
7	Hungary	1
8	Ireland	1
9	Latvia	1
10	Norway	1
11	Poland	1
12	Romania	1
13	Spain	6
14	Turkey	1
	Total	26

Table 1 – Countries contributing to e-FISCAL survey

Except for the 26 completed questionnaires reported above there are another 10 questionnaires that were not suitable for processing. Three out these 10 questionnaires corresponded to early attempts of filling in the questionnaire that were eventually completed and therefore included in the sample while the other 7 ended up having excessive missing information and were excluded for the subsequent analysis.

In order for a questionnaire to be considered as completed and suitable for further analysis, it should contain information corresponding to more than 75% of the following questionnaire categories.

Moreover, it should be noted that we have identified some answers that needed clarifications. We did not use any information that we considered to be out of range (mainly based on benchmarks) for the analysis that is subsequently presented¹⁷. Clarifications on responses were sought either through e-mail exchange or are planned to be discussed during the e-FISCAL summer workshop in July 2012.

Question Number	Short description
Descriptive	
2.2	Type of services
2.3	Type of institute
2.4	Country
2.5-2.7	Position - name - e-mail of responded (s)
2.8	Names of infrastructure (HPC and HTP) and the sites
Computing and hardware	
3.1	Logical CPU information
3.2	Configuration information
3.3	Wall clock time information
3.4	Storage TB information
3.5	Prices of logical CPUs and Storage
3.6	Depreciation information
3.7	Interconnect equipment cost information %
3.8	Support contract costs information%
Auxiliary equipment	
4.1	Auxiliary equipment cost information
4.2	Auxiliary equipment elements
Software costs	
5.1 or 5.2	Software costs information
Personnel costs	
6.1 or 6.2 & 6.3	Personnel cost information
Site information	
7.1	Site information (at least one indication for space/electricity consumption/ PUE/Computing to cooling consumption)
Connectivity costs	
8.3	Connectivity costs
Overhead costs	

¹⁷ For example we received a PUE value below 1 (apparently mixed with the reverse of PUE, i.e. Data Centre Infrastructure efficiency (DCiE)).

9.1	Overhead costs
Business Models and Funding	
10.1	Funding info
10.2-10.6	Business model related information and Green IT

Table 2 – Questionnaire overview- cost categories and other related information

The respondents in their vast majority (all but one) provide both computing and coordination activities as shown in Table 3.

What type of services does your institute provide?		
Answer Options	Response Percent	Response Count
Coordination (no resources)	3,8%	1
Computing services (CPU, storage, etc.)	46,2%	12
Both	50,0%	13
answered question		26

Table 3 – Types of services represented in e-FISCAL survey

Out of the 26 respondents, 12 indicated that their institute is only part of NGI/EGI, 3 reported participating into National HPC infrastructure/PRACE while another 10 participate in both (see Table 4).

To which e-infrastructure is your institute's infrastructure part of?		
Answer Options	Response Percent	Response Count
NGI/EGI	42,9%	12
National HPC infrastructure/PRACE	10,7%	3
Both	28,6%	8
Other (please specify)*	17,9%	5
answered question		26

* There are two questionnaires indicating that the sites are “National HPC infrastructure/PRACE and other” and “Both and other”

Table 4 – Institute's infrastructure participation represented in e-FISCAL survey

The respondents belong to institutes that have multiple roles ending up in a vast number of combinations. The majority of institutes participating in e-FISCAL survey (18 out of 26) are NGI Resource Centres (CPU, storage, etc.), while another 14 out of 26 are HPC centres. An analytical review of these combinations is presented in Table 5. From the table one can understand that there is no easy

separation between HTC and HPC centres, especially when comparing the high-end HTC ones with the low-end HPC ones. This was revealed during the edition of one of the intended publications of the project (for the eChallenges 2012 conference) and was also raised with the e-Infrastructure community during the e-IRG workshop. As there are several definitions of these two terms, some of which are not satisfactory (as the Wikipedia one), the project decided to raise the attention of the e-Infrastructure community and especially EGI and PRACE. In the Terms and Definitions section we have adopted the EGI definition.

Type of institute	Total
1,2,3	1
1,2,3,4	1
1,2,4	3
1,2,3,4,5	1
1,3	1
2	7
2,3,4	1
2,4	2
2,4,5	1
2,5	1
3,4	1
4	3
4,5	1
5	2
Total	26

Where:

1= NGI (coordinating body), 2= NGI Resource Centre (CPU, storage, etc.), 3= PRACE country coordinator, 4= HPC centre and 5 = Other (please specify)

Table 5 – Institute’s roles represented in e-FISCAL survey

3. LITERATURE REVIEW

We had – based on expertise in the consortium – a fairly large set of papers to start with. This was completed through usual desk research methods as well as publishing the state of the art repository (<http://www.efiscal.eu/state-of-the-art>) and attempting to contact authors of the papers included to see if they had additional suggestions. This engagement with the authors produced five additional references as well as an interesting additional tool for project outreach efforts. The desk research was aimed at finding relevant project reports, conference papers, academic papers (accessed through university subscriptions), Internet articles, professional reports, etc.

3.1 Overview of the state of the art

The first step in the literature review process is related to the identification of the proper methodology in assessing the e-infrastructure cost. Two methodologies stood out; the Total Cost of Ownership (TCO)¹⁸ methodology and the full costing methodology¹⁹. However as it has been presented in section 2.1 we instead of adopting either of those we developed a hybrid model that could satisfy the e-FISCAL project needs under the existing data and time constraints. The procedure followed has been presented in Del. 2.1. We briefly discuss later on the characteristics of TCO and full costing methodology for reasons of completeness. TCO paradigm is a useful tool when the cost of a specific project is to be assessed over its useful life. The developers of such models adopt a *forward looking* stance covering all *expected* costs over the project's lifetime. This cost is divided by the anticipated useful life to come up with a yearly cost. However, in order for this analysis to be precise enough, several details have to be taken into consideration as basis of the future cost estimates. Such details are very difficult to gather consistently across several organisations. Nevertheless, this procedure has been extensively used in literature (Nazir and Sorensen (2010), Walker (2009)).

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. However, the reliability of the cost data depends heavily on the robustness of the cost accounting system. And even with highly developed accounting practices, costs that are not funded by the organization are not registered as parts of the total costs.

After concluding with the selection- development of the costing model, the review of the state of the art has been conducted in other fronts. One category of papers and studies reviewed was dealing with the cost per core/hour under different settings. The ultimate goal is to compare in-house costs with commercial cloud offerings.

¹⁸ http://en.wikipedia.org/wiki/Total_cost_of_ownership.

¹⁹ http://en.wikipedia.org/wiki/Full_Cost_Accounting.

A comprehensive example on this subject is presented in the Magellan final report (2011). Magellan is a project funded through the U.S. Department of Energy (DOE) Office of Advanced Scientific Computing Research (ASCR). Its goal is to investigate the potential role of cloud computing in addressing the computing needs for the DOE Office of Science (SC). In the conclusions of the project's final report it is stated, among others, that "the cost analysis shows that DOE centres are cost competitive, typically 3-7x less expensive when compared to commercial cloud providers". Hawtin et al (2012) in a cost analysis of cloud computing for research (on behalf of EPSRC and JISC) report that the more powerful cloud computing instances, rented on an hourly basis, appear to be one-and-a-half to two times more expensive per core-hour than well-managed, locally-provided clusters in modern data centres operating at high utilisation levels. However, other purchasing models (such as 'Reserved Instances') can reduce the costs to parity.

Smith in his thesis (Smith, 2011) presents a model for calculating the base cost for a core-hour of computation in Purdue University's campus grid. With the cost model developed, the author analyses the benefits gained from using the grid, based on the number of hours of delivered, number of computations completed, and the number of users and faculty members served. In the same vein, a practical example of cost calculations and comparisons with cloud providers refers to Hyak shared compute facility at the University of Washington²⁰. Carlyle et al. (2010), present a case study of costs incurred by faculty end-users of Purdue University's HPC "community cluster" program. The authors develop and present a per node-hour cloud computing equivalent cost that is based upon actual usage patterns of the community cluster participants and is suitable for direct comparison to hourly costs charged by one commercial cloud computing provider. They find that the majority of community cluster participants incur substantially lower out-of-pocket costs in this community cluster program than in purchasing cloud computing HPC products. Another case study is found in Walker (2009). Walker proposes a modelling tool that can quantitatively compare the cost of leasing CPU time from online services to that of purchasing and using a server cluster of equivalent capability. A first attempt to approximate the cost of EGI is found in Cohen and Karagiannis (2011). In this study, conducted within the e-IRGSP2 project, an approximation of the cost of the EGI pan-European grid infrastructure for 2009 by extrapolating the cost of a few (seven) selected NGIs as example cases is attempted. Moreover, this study presents the cost per CPU core hour under alternative CPU utilization rates. As a final step rough comparisons between these costs and Amazon EC2 prices are made.

Risch and Altmann (2008) analyse the question whether using the Grid is financially advantageous, using the Amazon.com EC2 service as a reference. To perform this analysis they calculate the costs of computing resources in different usage scenarios, reflecting cases where Grid resources and in-house resources are used. The comparison of the costs reveals that while the Grid is cheaper in the short term, it is not a good investment in the long term and, thus, the existence of a Grid economy will not lead to an end of ownership but rather to a reduction of in-house resources and more efficient resource usage.

²⁰ <http://escience.washington.edu/get-help-now/hyak-operating-costs-and-comparison-commercial-alternatives>

Another significant part of literature analyses the aspects of Cloud computing. Lin and Chen (2012) investigated how cloud computing is understood by IT professionals and the concerns that IT professionals have in regard to the adoption of cloud services. The findings of this study conducted through interviews in Taiwan suggest that while the benefits of cloud computing such as its computational power and ability to help companies save costs are often mentioned in the literature, the primary concerns that IT managers and software engineers have are compatibility of the cloud with companies' policy, Information Systems (IS) development environment, and business needs. The findings also suggest that most IT companies would not adopt cloud computing until the uncertainties associated with cloud computing, e.g. security and standardisation are reduced and successful business models have emerged. Marston et al. (2011) discuss cloud computing from several fronts. They identify the strengths (e.g. reduced infrastructure costs and energy savings as well as reduced upgrades and maintenance costs), weaknesses (e.g. the loss of physical control of the data that is put on the cloud), opportunities (e.g. energy consumption and carbon footprint reduction) and threats (e.g. lack of standards that may force customers into locked, proprietary systems that will gradually cost more and more over time) for the cloud computing industry. They also discuss IS policy issues which cover a broad set of topics, from data privacy and data security to data ownership and audit. Hammond et al. (2010) discuss cloud computing for research in the areas of compute and storage in relation to Infrastructure as a Service (IaaS) and Platform as a Service but not Software as a Service. Authors argue that critical thinking is still required from researchers and institutions as to what data storage or compute solution is most appropriate given functional requirements, budget, security, reliability, trust, etc. as well as the cloud services currently on offer. The EGI (2011) report aims to evaluate technologies such "Infrastructure as a Service", "Platform as a Service" and "Software as a Service", understand how they relate to EGI, and build a foundation for the integration of cloud and virtualisation²¹ into the European production infrastructure. The report includes a cost analysis and comparisons to current market offers. The EGEE (2008) report lies within the same realm. This report compares grid and cloud computing services, taking a practical look at implementations of the Enabling Grids for E-science (EGEE) project for grid and the Amazon Web Service (AWS) for cloud. Taking performance, scale, ease of use, costs, functionality and other aspects into consideration, the report looks at the overall opportunity that converging cloud and grid services can bring to users. Misra and Mondel (2011) on the other hand propose a model to assist companies by analysing several characteristics of their own business as well as pre-existing IT resources to identify their favourability in the migration to the Cloud Architecture. They also develop a general Return on Investment (ROI) model that embodies various intangible impacts of Cloud Computing, apart from the cost, in order to provide a tool for a broader perspective and insight cloud computing decision assessment.

Another research area includes papers dealing with electricity cost and cost of data centres in general. The electricity efficiency issues fall within the broader area of Green IT. In Koomey et al. (2009) the relationship between the processing power of computers and the electricity required to deliver that performance is analysed. Rasmussen (2011) proposes a method for measuring total cost of ownership

²¹ EGI (2012) discusses in more depth EGI.eu strategy towards virtualization.

(TCO) of data centre network room physical infrastructure and relates these costs to the overall information technology infrastructure with examples. In several cases PUE²² has become a measure of datacentre efficiency. The Cordis (2011) report discusses potential actions that could contribute towards improving efficiency in data centres and reducing costs. Such actions include among other new approaches in minimizing cooling requirements, improvements to uninterruptible power supplies (UPSs) and usage of renewable energy sources. According to a recent study in UK comprising 27 UK colleges and universities (Hopkinson and James, 2012), the total energy for all PCs (including desktops, laptops, thin clients and monitors) ranged from 144 kWh/y to 587 kWh/year. Especially in HPC-intensive universities the average energy consumption per PC is 309 kWh/y.

The pricing of cloud services is also another issue. The e-Science Institute²³ focuses on the steady decrease in prices offered by Amazon for their services. They comment that over time, the price to rent 1 unit of resources for three years of continuous usage has fallen dramatically as Amazon offered new instance types, offered new long-term pricing plans, and lowered prices outright across the board.

3.2 Conclusions retrieved by the state of the art review

It is evident from the analysis described above that the discussion of comparing cloud services prices with in-house costs has been placed in the focal point of discussions. Cost issues along with confidentiality and other policy considerations play an important role in the decision process that relates to cloud migration. Moreover, cost issues are of high priority relating to sustainability assessment and Green IT evaluations. As for the empirical examples found in current literature, the cost per logical CPU/hour is not homogenous. On the contrary, it is reliant on the “production conditions”. Economies of scale and usage intensiveness heavily affect cost per logical CPU /hour. The table below summarizes the results of the literature review.

Reference	Cost per core hour	Comments
Hawtin et al. (2012)	£0.05 - £0.07	Study for JISC UK - Differences between institutions reviewed
Cohen and Karagiannis (2011)	€ 0.0854 – € 0.1356	Stratified sample of EGI centres - Assuming 60% utilization ratio – storage cost included
Cohen and Karagiannis (2011)	€ 0.0782 – € 0.1020	Stratified sample of EGI centres - Assuming 60% utilization ratio – storage cost excluded
US DoE - Magellan report (2011)	\$ 0.018	Hopper system – National Energy Research Scientific Computing Centre- including storage sub- system

²² Power Usage efficiency (PUE) is determined by dividing the amount of power entering a datacentre by the power to use to run the computer infrastructure within it.

²³ <http://escience.washington.edu/blog/cloud-economics-visualizing-aws-prices-over-time>

Smith (2011)	\$ 0.039	Purdue campus, USA
University of Washington	\$ 0.051	Hyak cluster, USA

Table 6 – Literature review summary in relation to Cost per logical CPU/hour or Cost per core/hour

4. DATA ANALYSIS

4.1 Cost analysis – Presenting the input values

The first part of the analysis is devoted to the summary presentation of the answers on the questionnaire questions relating to costing. We go through the questions and provide some descriptive statistics. These statistics present the range (minimum values vs. maximum values) of observations as well as the average and the median²⁴. In parallel we discuss these statistics with the input retrieved through the EGI compendium.

It has to be noted that respondents in several cases mention that their answers are based on approximations. This is not surprising and respondents provide very concrete arguments to support the reason why this is happening. Firstly, as several answers are given for multiple sites the numbers provided correspond to an average estimation per site. This is especially true for the operating costs that are reliant on the operating environments in which energy/cost efficiencies differ. Secondly, several contributions relating to hosting/operations in the sites are provided in kind from partners and therefore they are neither registered nor directly visible. Although this dimension makes the total cost estimation difficult and the total cost structure rather complex, cautious conclusions can be derived.

Logical CPU information

The logical CPU information of the sample is presented in Table 7.

Please present information in relation to the total number of “logical” CPUs (i.e. number of processing cores) of the NGI site/ HPC Centre available at the end of years 2010 and 2011.				
Answer Options	Min	Max	Average	Median
Logical CPUs as on 31/12/2010	48	19100	3592	1048
Logical CPUs as on 31/12/2011	72	23100	5072	2586
answered question				26

Table 7 – Logical CPU information

Storage information

Information about the Storage capacity of the sample is presented in Table 8.

Please present information in relation to the TB of storage devices available in the NGI site/ HPC Centre at the end of year 2010 and at the end of year 2011.				
Answer Options	Min	Max	Average	Median
Disk Storage in TB as on 31/12/2010	1	2493	406	139

²⁴ The median is described as the numeric value separating the higher half of a sample, a population, or a probability distribution, from the lower half.

Disk Storage in TB as on 31/12/2011	3	3308	657	333
Tape Storage in TB as on 31/12/2010	0	2528	377	0
Tape Storage in TB as on 31/12/2011	0	5176	592	0
answered question			26	

Table 8 – Storage information

Comparing the results in Table 7 and 8 with EGI compendium²⁵ input it seems that our median respondent is by 22% larger in terms of logical CPUs as of 31/12/2011 than EGI compendium respondents (median 2,586 logical CPUs to median 2,126 logical CPUs, respectively). There are no material differences in terms of the median disk storage (median 333 TB vs. median 328 TB) and tape storage (median 0 TB to median 0 TB).

Cost information

Information about the acquisition costs of both hardware and storage of the sample is presented in Table 9.

Please present the average acquisition (i.e. purchase) cost per logical CPU and the average cost per TB acquisition in 2010 and 2011. In case you have no data for 2011 please use approximations based on the most recent procurements or budget data. Note: Please do not include any hardware support contract costs or software costs in the values presented below					
Answer Options	Min	Max	Average	Median	Answered questions
Cost per logical CPU in € in 2010	100	3150	460	300	17
Cost per logical CPU in € in 2011	80	3400	439	225	20
Cost per TB/ Tapes in € in 2010	50	150	97	94	4
Cost per TB/ Tapes in € in 2011	37	125	79	78	4
Cost per TB/ Disks in € in 2010	65	6000	704	315	15
Cost per TB/ Disks in € in 2011	80	3000	503	250	15

Table 9 – Acquisition cost of logical CPU and storage information

The cost per logical CPU and the cost per TB exhibit significant differences. As one respondent commented these costs are rather dependent on the choice of compute/storage technology. It can be easily seen that the number of answers to this question are less than the total sample. According to informal discussions the disclosure of procurement information related to prices is considered sensitive data and therefore some respondents refrained from presenting their case²⁶. The limited number of responses in relation to cost per TB for tapes is probably related to the fact that not all respondents use tapes as a storage medium.

²⁵ By excluding CERN and two very small NGIs with 16 logical CPUs in total.

²⁶ We expanded our information set regarding hardware (computing and storage) procurement prices by getting information from the EGI compendium study. This type of information was used for cross-checking purposes. The median cost per logical CPU in EGI compendium was € 192 (6 responses) while the cost per TB € 373.

Depreciation rates

The average useful lives for hardware and software are presented in Table 10. Depreciation rate can be calculated accordingly.

Please indicate the period in number of years that corresponds to the average useful economic life (depreciation period) of the following assets according to the policy followed by the NGI site/ HPC Centre.					
Answer Options	Min	Max	Average	Median	Answered questions
Average useful life in years for CPUs	3	10	5	5	23
Average useful life in years for tape storage devices	3	12	7	5	12
Average useful life in years for disk storage devices	3	20	6	5	23

Table 10 – Useful life information

Analysing Table 10 allows us to make an interesting observation. Typically the Total Cost of Ownership (TCO) exercises use primarily a four year duration (e.g. Walker (2009) and Magellan report (2011) and sometimes a three year duration (e.g. Nazir and Sørensen (2010)) – as the useful lifetime of a CPU. In our sample the median useful life to calculate annual depreciation is five years. This has a non-negligible impact on the capital costs accounted each year in order to form the total yearly cost. Nevertheless, there are instances of long depreciation periods which will require deeper study to understand the underlying conditions that permit such an treatment.

Cost relations over investment

In order to assess several other cost parameters that are related to CAPEX but at the same time are not easily identifiable and measurable, we asked respondents to provide indications on their relative size compared to the computing and storage investment. As these estimations were expected to be rough, respondents were given the flexibility instead of identifying a sole number (in our case percentage) to indicate a range of values (e.g. between 10%-20%). Therefore the following numbers are considered as high-level approximations and the assessment of their accuracy and robustness under different computing environments will be studied thoroughly in the next steps of the project.

Estimated cost relations of several parameters on computing and hardware storage				
	Min	Max	Average	Median
Please present an overall estimation of the related interconnect equipment costs (network devices, cables, etc.) as a percentage of the hardware acquisition cost*	0%	30%	10%	10%
Please present an overall estimation of the support contract costs (e.g. next-business-day hardware support costs) as a percentage of the hardware (CPUs and storage devices) acquisition cost **	0%	25%	6%	3%
If you were to equip the existing NGI site/ HPC Centre now	5%	35%	17%	20%

what would be the investment cost of all auxiliary equipment as percentage of the cost of acquiring computing and hardware storage capacity***				
Please make an estimation of the total cost of the related software (e.g. operating system, fabric layer / file system software (e.g. LSF, GPFS), software support contract costs, applications cost, 3rd party software cost, compilers, etc.) as a percentage of the hardware (CPUs and storage devices) acquisition cost	0%	15%	4%	2%

*A respondent commented that for CPU, first level switches are included in the servers' price. For instance, blade centres already include switches.

** Four respondents made comments on the "other" option in this question. Their comments could be summarised as follows: These support terms are included in hardware tender specifications (for 3 years or 3-5 years) and are difficult to separate (not typically included in contract as a line item). Apart from that some centres pay extra money for support while others don't.

*** Due to the wording of the question the analyses that follows has used 50% of the average percentage value of auxiliary equipment over investment in order to calculate the yearly cost per logical and per logical CPU/hour.

Table 11 – Cost relations over investment

The software costs are most accurately characterised as an operating expense in today's e-Infrastructure ecosystem. The majority of the system software used today is produced in collaborations (in which many of the service providers participate) and released under open source license, which can explain both the relatively low proportion of the costs and large variations between centres.

This contrast with the situation in many of the ICT environments outside the research domain, where the software licenses for applications such as enterprise resource planning (ERP) systems may form the biggest part of the capital expenditure, and a division of the software costs between CAPEX (license fees) and OPEX (service contracts) is necessary.

It should be noted that assessing the overall costs of the European e-Science ecosystem are outside of the scope of this study. However, it is possible to assume that algorithmic developments and software maintenance tasks that are directly related to research would require similar amounts of effort in dedicated or on-demand infrastructure. The overall structure and behaviour of the scientific software ecosystem is a topic of an e-IRG Task Force on Scientific Software that is due to release its final report later this year.

Auxiliary equipment information

All sites are well equipped in terms of auxiliary equipment. Air cooling is more popular than liquid cooling. While proceeding in the following months with the analysis of our data and the conduction of interviews and case studies, we will try to identify whether the type of cooling affects energy consumption or other parameters of operating expenses.

Please identify whether the NGI site/ HPC Centre possesses the following infrastructure elements.		
Answer Options	Response Percent	Response Count
UPS	96.2%	25
Air Cooling	88.5%	23
Liquid Cooling	57.7%	15
Power Generators	50.0%	13
Power Transformers	38.5%	10
Other (please specify)*		3

* Respondents commented in two cases that cooling equipment, power generators, power transformers etc. are part of the building, owned by a third party (e.g. of the university).

Moreover, another category identified by one respondent is fire detection and extinction equipment.

Table 12 – Auxiliary equipment information

Personnel information

Salary levels exhibit significant differences among countries (ref. Min and Max values in Table 13). It could be argued that personnel cost is a cost element that is highly dependent on local - country conditions; a phenomenon which is less evident as far as technology, type of infrastructure or architecture are concerned (i.e. the procurement of hardware or storage in two different counties is not expected to deviate to such an extend only due to the geographical area). Therefore the structure of our sample (i.e. counties participating in the survey) form some metrics of average and median personnel costs that for some other countries could be exceptionally high or low respectively. Interestingly enough the median cost per FTE from the EGI compendium data set is € 44.000 while the average cost is €46,200 very close to our findings.

Please provide the following information related to the cost of the personnel for 2010 and 2011 as well as an average yearly salary per FTE.				
Answer Options	Min	Max	Average	Median
Average yearly salary cost per FTE (gross salary plus employee benefits and bonuses) in '000 € in 2010	15	103	48.55	47.60
Average yearly salary cost per FTE (gross salary plus employee benefits and bonuses) in '000 € in 2011	15	103	49.31	47.60

Table 13 – Salary information

Space and PUE information

In the following table information about the site space and the Power Usage Effectiveness (PUE) are discussed. We will revisit with respondents the values of PUE as well as other metrics asked within this set of questions relating to operating costs. We acknowledge that in several cases this type of information is not systematically kept. However, due to the increasing awareness on Green IT aspects these data would gain more and more importance especially due to the increased cost of electricity consumption.

Please fill in the following information related to the cost and operating characteristics of the NGI site/ HPC Centre for 2010 and 2011.				
Answer Options	Min	Max	Average	Median
Site centre space in m2 (2010)	3	1,000	145	88
Site centre space in m2 (2011)	6	1,000	163	113
Power Usage Effectiveness in 2010	1.4	2.50	1.68	1.51
Power Usage Effectiveness in 2011	1.4	2.24	1.65	1.50

Table 14 – Site information

The questions about network connectivity costs (question 8.3) and questions relating to overhead costs were only sparsely completed (questions 9.1 – 9.2). As for the connectivity, almost half of the respondents (12 out of 26) do not pay for connectivity to Internet/NREN while the other half (11 out of 26) do. Moreover, when these pieces of information were reported the magnitude of costs for both connectivity and other overhead was in general small and it looks like, at least for the current sample, that it would not make a difference in the calculated amounts. The overhead costs refer to travelling expenses, conferences, training and insurance. Nevertheless, we could come back to this issue during interviews and more thorough discussions with the financial experts of our sample respondents.

As there is not enough information we do not report any values about these two cost categories. However, we keep these cost categories as part of the overall cost breakdown with zero amounts. Notwithstanding that their non-inclusion is not expected to have a material effect, we accept that by taking into consideration such costs, cost calculations would have been driven, marginally, up.

Cost analysis – Metrics calculation

The information above provides adequate data in order to calculate, some preliminary, relevant metrics in order to estimate the cost per logical CPU and the cost per logical CPU/hour. These data would permit the identification of an average yearly cost per logical CPU and an average cost per logical CPU/ hour that correspond to the sample. Cost per logical CPU/hour is considered to be one of the more direct methods to compare costs with cloud offerings (Magellan report, 2011). However, as it has been shown in the tables above the sample is not homogenous. This is evident by the deviation reported in terms of range values (minimum and maximum values) compared to the average or median ones.

As this report is focused on presenting the preliminary findings of our project we decided to aggregate all data in a logical CPU level and make our unit of analysis the logical CPU. We are aware that this procedure does not adequately account for the existence of economies of scale and does not correspond to the specific circumstances of each and every participant in the study. However, as we both employ information for average and median values we expect to make a fair approximation of the “average” case. Additionally, the logical CPU unit of analysis permits the presentation of results in a way that is easy to understand and follow. It is a therefore a benchmark that could be used for high-level comparisons.

A graphical representation of the methodology followed (as discussed in section 2.1) is given in Figure 3.

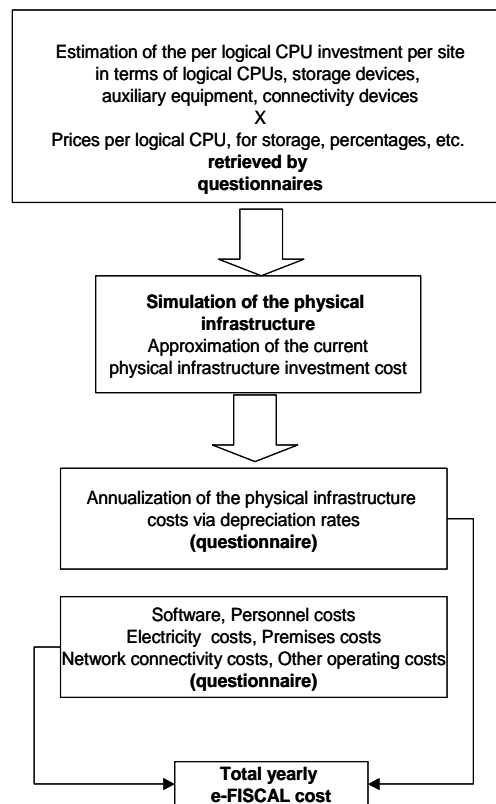


Figure 4: Graphical representation of e-FISCAL methodology

By applying the methodology depicted in Figure 4 the following relations were revealed.

	2010		2011	
	Average	Median	Average	Median
Average Cost per logical CPU	460	300	439	225
Average cost per Tape storage TB	97	94	79	78
Average cost per Disk storage TB	704	315	503	250
Primary investment (CPUs + storage)/ Logical CPU*	611.1	321.0	550.3	242.1
Other investment (interconnection + software contracts+ auxiliary)**	148.21	67.40	133.45	50.84
Overall invested capital / logical CPU	759.35	388.36	683.73	292.93
Total yearly CAPEX/ Logical CPU ***	149.0	77.7	134.7	58.6
Total yearly operating costs (OPEX)/ Logical CPU ****	386.9	180.3	288.0	138.5
Total yearly cost/ Logical CPU	535.9	258.0	422.7	197.1
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%

*To calculate the primary investment, the relation between logical CPUs and storage capacity in TB per storage medium (disk or tape) was used

** To calculate the cost of other investments, percentage relations identified through questionnaires were used

***The yearly investment cost derived after applying depreciation rates.

**** The analysis of operating costs is presented in Table 16

Table 15 – Yearly Cost per logical CPU (in €)

The acquisition costs of computing and storage devices shows a decreasing trend between years 2010 and 2011 (consistent with expectations), while investments that relate to interconnection devices, auxiliary equipment etc. form a non-negligible part of the investment cost. The total yearly cost per logical CPU ranges from €197.10/year (2011, median) to €535.90 (2010, average).

In Table 16 the average numbers of FTEs per logical CPU²⁷ and the m2 per logical CPU exhibits a decreasing trend from 2010 to 2011 as well. This trend is consistent with economies of scale. More specifically, this may be largely due to the impact of new multi-core servers: same number of physical servers administered by the same teams, but with higher number of cores. Similarly, the electricity consumption per logical CPU shows a decreasing trend. According to (Koomey, 2011) servers in 2010 (compared to 2005) have much higher processing power, more memory, faster network connections, more components and larger power supplies. However, they have improved power management and other improvements that reduce electricity consumption. As for the electricity cost, an average €0.10/kwh²⁸ has been used while for premises costs an average of €100/m2 per year is applied²⁹. Both inputs are conservative..

²⁷ On the basis of the responses on EGI compendium the median FTEs/1000 CPUs is 2.02 which is very close to our findings while the average value is 13.22 per FTEs/1000 CPUs. This significant difference is due to the step-wise cost behaviour of personnel costs as human effort cannot be acquired proportionally to the demand.

²⁸ <http://www.energy.eu/#industrial>. End-user energy prices for industrial consumers reference November 2011

²⁹ We have used as a reference prices indicated on <http://www.globalpropertyguide.com/faq/guide-sqm-prices->

	2010		2011	
	Average	Median	Average	Median
Software value	23.39	4.81	21.06	3.63
FTEs/1000 CPUs	6.54	2.80	4.71	2.08
Average cost per FTE	48,547	47,600	49,305	47.600
Personnel cost / Logical CPU	317.5	133.2	232.4	99.2
m2/'1000 Logical CPU	80.93	59.58	60.38	49.34
Average cost per m2	100	100	100	100
Site premises cost	8.093	5.958	6.038	4.934
Electricity consumption per logical CPU in kWh	379.41	363.41	284.80	307.74
Cost /kWh	0.1	0.1	0,1	0.1
Electricity cost	37.941	36.341	28.480	30.774
Connectivity costs ³⁰	0	0	0	0
Other costs ¹²	0	0	0	0
Total operating costs	386.9	180.3	288.0	138.5

Table 16 – Operating Yearly Cost per logical CPU (in €)

4.2 Calculation of the logical CPU cost/hour

The following Table (Table 17) presents the cost per logical CPU/hour.

	2010		2011	
	Average	Median	Average	Median
Total yearly cost/ Logical CPU	535.9	258.0	422.7	197.1
Yearly logical CPU minutes	8760	8760	8760	8760
Utilization rate*	59%	59%	58%	67%
Cost per logical CPU/hour	0.1036	0.0499	0.0837	0.0337

* In the questionnaire we have asked for logical wall clock time in order to assess utilization rate³¹.

Table 17 – Cost per logical CPU/hour (in €)

The utilization rate in the e-FISCAL project has been calculated by taking into account yearly logical wall clock time and available logical CPUs at the end of the period (either 31/12/2010 or 31/12/2011). This assumes that the number of logical CPUs reported at the end of the period is available throughout the year. If this is not the case (e.g. because of an infrastructure upgrade towards the end of the year), the utilization rate calculated is underestimated.

[rents-yields](#). Nevertheless, the €100/m2 per year is a rough estimation.

³⁰ To be revisited in the second year of the project

³¹ This refers to a mixture of EGI, PRACE and other sites not integrated in these e-Infrastructures. As an example, for 2011, EGI reports an utilisation rate of 71,3%

The cost per logical CPU hour are close to the ones discussed in Hawtin et al. (2012) that concluded that the cluster costs come in the region of 0.05€ - 0.07€ per core hour (based on discussions with research computing managers and information shared in confidence), but with significant variations between organisations. The costs in Table 17 are also in the same magnitude with cost/Logical CPU in Cohen and Karagiannis (2011). The above reported costs in the US are in some cases lower, as they are more centralised and thus lack the higher personnel costs of the distributed European infrastructures (that is the case for EGI and PRACE centres)³²: the Hopper system operated by National Energy Research Scientific Computing Centre (Magellan Report, 2011) the reported cost is of \$0.018/core hour, including some storage related costs. According to Smith (2011) the total cost of a core-hour in the Purdue campus grid is \$0.039847 while in the Hyak cluster of the University of Washington³³ the reported yearly cost per node is \$2,794 and therefore the cost per core hour is \$ 0.051358. Thus the results of the survey are in broad agreement with other published results. The above values have been summarized in Table 6.

In the following tables a breakdown of the total yearly cost per logical CPU into broad cost categories is presented. A reason why the contribution of OPEX compared to CAPEX is significant is partially due to the rather high average depreciation period used in the analysis. The application of a shorter useful life would increase the yearly cost per logical CPU and at the same time the percentage contribution of CAPEX to the total. Such considerations are explored in next section.

Cost break down / logical CPU	2010		2010	
	Average in €	Median in €	Average %	Median %
Depreciation Logical CPUs	92,9	60,0	17%	23%
Depreciation storage	26,2	4,2	5%	2%
Depreciation other	29,9	13,5	6%	5%
Software	23,39	4,81	4%	2%
Personnel	317,50	133,21	59%	52%
Premises cost	8,09	5,96	2%	2%
Electricity cost	37,94	36,34	7%	14%
Other cost	0,00	0,00	0%	0%
Total yearly cost	535,9	258,0	100%	100%

Table 18 – Cost per logical CPU per cost category (in €) and % for 2010

Cost break down / logical CPU	2011		2011	
	Average in €	Median in €	Average %	Median %
Depreciation Logical CPUs	88,6	44,9	21%	23%
Depreciation storage	19,2	3,5	5%	2%
Depreciation other	26,9	10,2	6%	5%

³² It is obvious that the personnel costs for a very large computing centre (in the order of 100.000 cores) can show economy of scale if compared to the same capacity distributed in smaller sites that are federated together. Further analysis will explore this issue in more detail.

³³ <http://escience.washington.edu/get-help-now/hyak-operating-costs-and-comparison-commercial-alternatives>

Software	21,06	3,63	5%	2%
Personnel	232,39	99,17	55%	50%
Premises cost	6,04	4,93	1%	3%
Electricity cost	28,48	30,77	7%	16%
Other cost	0,00	0,00	0%	0%
Total yearly cost	422,7	197,1	100%	100%

Table 19 – Cost per logical CPU per cost category (in €) and % for 2011

4.3 Sensitivity analysis

Several parameters may affect the logical CPU/hour cost. A summary of these factors follows:

- **Utilization rates:**
 - The utilization rate (number of CPU hours actually used compared to the theoretical maximum) plays an important role in calculating the cost per logical CPU/hour. CPU hours that are not used could be considered as are wasted resources whose cost is added to the hours been used (Hawtin et al., 2012).
- **Deprecation rates:**
 - The prolongation of hardware and auxiliary equipment useful life ends up in decreasing the yearly cost per logical CPU and the corresponding cost per logical CPU/hour.
- **Auxiliary equipment:**
 - In the study we took a modest stance towards auxiliary equipment as a part of the infrastructure in several cases is shared. Assuming that no shared resources of this type exist would increase both yearly cost per logical CPU and the corresponding cost per logical CPU/hour.
- **Salaries and premises costs:**
 - These costs have a very local focus and are highly influenced by national circumstances.
- **Electricity effectiveness and cooling**
 - The PUE of the respondents as well as the climate conditions that prevail in the site area affect electricity consumption which in turn affects OPEX calculation.

In the table below, the results of a sensitivity analysis that corresponds to the first three factors treated independently and reflect adjustments to the basic scenario are presented.

	2010		2011	
	Average	Median	Average	Median
Scenario 1: Utilization rate of 80%	0.0765	0.0368	0.0603	0.0281
Scenario 2: 3-year CPU and other investment useful life	0.0987	0.0400	0.1178	0.0588
Scenario 3: Auxiliary equipment costs as in Table 12	0.1057	0.0514	0.0858	0.0347
Basic Scenario	0.1036	0.0499	0.0837	0.0337

Table 20 – Cost per logical CPU / hour (in €) under alternative scenarios

4.4 Business models

The last part of the questionnaire contained questions regarding cloud computing, Green IT aspects and sustainability considerations. The time horizon of these questions tried to balance recent past actions (i.e. actions that have been realized in 2011) and future prospects.

As it can be observed from Table 21 the majority of respondents are not allowed to use their funding neither project funding (82%) nor not – project funding (73%) to buy cloud related services.

Are you allowed to use your funding to buy Cloud related services?					
Answer Options	Yes	Yes (%)	No	No (%)	Response Count
Project funding (matching funds included)	4	18%	18	82%	22
Non project funding (e.g. national budget subsidies)	6	27%	16	73%	22

Table 21 – Allowed use of founding for buying cloud computing services

This has possibly contributed into limited usage of cloud computing services during 2011. As it is apparent from Table 22 the most frequently used cloud computing service is Software as a Service (SaaS), which counts 9 answers followed by Infrastructure as a Service (IaaS) with 7 answers.

Please answer the following questions in relation to the use of cloud computing in 2011: Note: "Use" means free of charge		
Answer Options	Use	Buy
Did you use (buy) Infrastructure as a Service (e.g. Amazon EC2) in 2011?	4	3
Did you use (buy) Platform as a Service (e.g. Microsoft Azure) in 2011?	2	0
Did you use (buy) Software as a Service (e.g. Google Docs, Microsoft Live services) in 2011?	7	2
Did you use (buy) disk storage services from external providers in 2011?	2	0
Did you use (buy) tape storage services from external providers in 2011?	0	0

Table 22 – Using and buying cloud computing services in 2011

However the prospects for the future are in favour of a more intensive, compared to 2011, cloud computing services usage (Table 23). Nevertheless, most respondents are more prone to use (free of charge) cloud computing services compared to those expressing their intention to buy such services.

Please answer the following questions in relation to the use of cloud computing in the future:		
Answer Options	Use	Buy
Do you intend to use (buy) Infrastructure as a Service in the future?	8	5
Do you intend to use (buy) Platform as a Service in the future?	6	0
Do you intend to use (buy) Software as a Service in the future?	10	3
Do you intend to use (buy) disk storage services from external providers in the future?	9	1
Do you intend to use (buy) tape storage services from external providers in the future?	3	0

future?		
---------	--	--

Table 23 – Using and buying cloud computing services from 2012 onwards

Moreover, one respondent raised some concerns regarding cloud computing that refer to three dimensions: a) the cost of using cloud services to transfer and store large amounts of data (e.g. hundreds of TB), b) security issues related to the safeguarding of data confidentiality (i.e. confidential data stored to third party storage devices), and c) CPU capacity (that is translated into additional costs) necessary to encrypt and regularly decrypt the data stored in third party premises.

Respondents proved very positive towards Green IT initiatives and some of them had already embarked into projects towards this direction³⁴. Energy efficiency considerations had played a role in investment decision of almost all respondents (21 out of 23 answers) while changes aiming at increasing energy efficiency were applied in more than half the respondents (12 out of 21). Finally, eight out of twenty-two respondents reported that in 2011 had recycled CPUs and storage devices.

Please answer the following questions in relation to "Green IT"		
Answer Options	Yes	No
Did you recycle CPUs or storage devices during 2011?	8	14
Do you plan to use some sort of "Green IT" in the future?	22	2
Had energy efficiency considerations influenced your acquisition decisions in 2011?	21	2
Did you make any changes in your hardware/software environment to increase the energy efficiency in 2011?	12	9
Did you use any form of "Green IT" in 2011?	12	8

Table 24 – Green IT related actions in 2011 and future prospects

Respondents generously provided details on their plans and experiences with Green IT. The majority proved very positive towards Green IT initiatives and some of them had already embarked into projects towards this direction. Examples of Green IT initiatives reported by respondents are summarized as follows:

- Buying energy efficient servers (improve performance per Watt).
- Virtualising more IT services.
- Reusing heat from servers to warm water for nearby buildings.
- Buying new hardware to replace old hardware.
- Building new datacentres.
- Applying efficient cooling systems.
- Exploitation of external temperature in order to use free cooling, fully or partially, during the whole.
- Machine rooms in the national infrastructure capture/recycle heat from the compute systems.
- Reallocation of HPC systems. Such decisions are partially driven by Green IT considerations (no use of fossil energy sources) as well as by cost considerations (low electricity cost).

³⁴ It should be noted that as reported in Cordis (2012) report the European Commission is highly interested in new approaches to minimising cooling requirements, to reducing power consumption and increasing the re-use of heat generated.

- Improvement on airflow management by eliminating mixing and recirculation of hot equipment exhaust air and maximizing the return of air temperature to the computer room air-conditioning unit.
- Implementation of environment monitoring systems in order to better measure resources.

Finally as it is apparent in the following table the idea of imposing usage fees is not out of the question for several respondents that have included such provision in both the short and long term planning.

Please answer the following sustainability related questions:		
Answer Options	Yes	No
Do you have a short-term (e.g. 1–3 years) capacity and business plan for your computing infrastructure?	20	2
Do you have a long-term (e.g. 3–5 years) capacity and business plan for your computing infrastructure?	13	8
Is there a provision of any kind of usage fees in the short-term plan?	10	12
Is there a provision of any kind of usage fees in the long-term plan?	9	13

Table 25 – Sustainability related information

5. BENCHMARKING

This section details information about the benchmark study that is carried out in conjunction with the cost analysis for different computing infrastructures including High Performance, Throughput and Cloud Computing paradigms. Section 5.1 describes background and rationale for the benchmarking exercise. Sections 5.2 and 5.3 detail the system benchmarks and specification of the systems chosen for benchmarking. Section 5.3 summarise the preliminary findings of the benchmark study and finally Section 5.5 concludes with a summary and future directions.

5.1 Background and Motivation

A variety of computing paradigms such as the High Performance Computing, High Throughput Computing and Cloud computing are addressing the growing computing needs of the scientific and business community. In addition to the technical differences among these computing paradigms, the underlying business processes and models pose a big challenge while comparing different computing infrastructures. As the core focus of the e-FISCAL project is the cost and business models analysis of these computing platforms, it is useful to put the results of the e-FISCAL analysis in the context of solid technical performance data. Therefore, in the context of a typical research setting, the performance measures of the HPC, HTC and Cloud systems are needed.

System benchmarking has been a widely accepted and used approach to compare and evaluate the performance metrics for the computer systems. These performance metrics either focus on measuring the capability or capacity of the HPC or HTC systems respectively. The information provided by the benchmark studies will serve as an input to the cost model that is based on the data collected for the e-FISCAL questionnaire. The benchmarking results also help in verifying the developed cost model as mentioned earlier.

In the next section, we will detail the system benchmarks selected to evaluate various computing infrastructures (i.e. HPC, HTC and Cloud).

5.2 System Benchmarks

The performance or capability of the HPC systems is generally measured in FLOPS (i.e. Floating Point Operations Per Second). More specifically, it is also required to consider the computational and data movement characteristics of the large-scale parallel applications e.g. in the Molecular Dynamics (MD) domain. For the HTC system, the benchmarks are typically designed to scale with the performance of the HTC applications (e.g. High Energy Physics - HEP) on similar machines. The HTC benchmarks assess the performance measure of different configurations and detect latent bottlenecks and problems. The Cloud systems can be classified in one of the above categories, namely as HPC or HTC Clouds. Thus, the benchmarks for the HPC and HTC systems are applicable to the HPC and HTC Cloud systems respectively. Based on the above discussion, the following benchmark suites are selected to compare the HPC, HTC and Cloud systems: The NAS Parallel Benchmark³⁵ and the HEP-SPEC benchmark³⁶.

³⁵ <http://www.nas.nasa.gov/publications/npb.html>

³⁶ <http://hepix.caspur.it/benchmarks/doku.php>

The NAS Parallel Benchmark (NPB)

The NPB benchmark suite is a set of programs designed to evaluate the performance of HPC systems. The benchmark suite is comprised of eight programs: i.e. five parallel kernels and three simulated applications. The NPB suite expresses the computation and data movement involved in the typical scientific applications e.g. Molecular Dynamics applications. The benchmark suite comes with different classes, each representing a different problem size. For the benchmarking in the context of this deliverable, we have selected the “Class B” problem size for various programs (i.e. BT, SP, EP, MG, CG, IS, LU and FT). The Class B programs characteristic are summarised in Table below:

Program	Description	Problem size	Memory (Mw)
EP	Monte Carlo kernel to compute the solution of an integral – Embarrassingly parallel	2^{30}	18
MG	Multi-grid kernel to compute the solution of the 3D Poisson equation	256^3	59
CG	Kernel to compute the smallest eigenvalue of a symmetric positive definite matrix	75000	97
FT	Kernel to solve a 3D partial difference equation using an FFT based method	512x256x256	162
IS	Parallel sort kernel based on bucket sort	2^{25}	114
LU	Computational Fluid Dynamics (CFD) application using symmetric successive over relaxation	102^3	122
SP	CFD application using the Beam-Warming approximate factorisation method	102^3	22
BT	CFD application using an implicit solution method	102^3	96

Table 26 – Class B Program Characteristics

In addition to different Classes for various problem sizes, the NPB benchmark suite supports a number of different parallelising technologies: i.e. OpenMP, MP, HPF and Java. In the context of this deliverable, we have compared the OpenMP and MPI versions of the NPB for the HPC and Cloud systems.

The HEP-SPEC benchmark

The High Energy Physics (HEP) HEP-SPEC benchmark is based on the SPEC (Standard Performance Evaluation Corporation) benchmark and is widely adopted by the HTC community for measuring the CPU performance on similar machines. The SPEC benchmark is a set of test applications which stress the processor with operations and algorithms used commonly in applications from the physics community, and provide the SPEC result used for example in describing the resources relevant to HEP applications, provided to the Grid infrastructure. Alternatively the HEP-SPEC benchmark suite can be used to assess the condition of different configurations and detect latent bottlenecks and problems.

Currently, we are working to define a test case for comparing the HTC infrastructures against the HTC equivalent of the Cloud offerings. Formulating such a test case will take place during summer 2012.

5.3 System Specification

In the benchmarking exercise, we have compared the Stokes HPC compute nodes at the Irish Centre for High End Computing (ICHEC) against the Amazon EC2 compute cluster instances. Below are the system specifications in more detail:

Stokes HPC System

Hardware specification

The Stokes HPC system is an SGI Altix ICE 8200EX cluster with 320 compute nodes. Each compute node has two Intel (Westmere) Xeon E5650 hex-core processors and 24GB of RAM. This results in a total of 3840 cores and 7680GB of RAM available for jobs. The nodes are interconnected via two planes of ConnectX Infiniband (DDR) providing high bandwidth and low latency for both computational communications and storage access. Storage is provided via a Panasas ActiveStor 5200 cluster with 143TB (formatted) of capacity to the compute nodes via a PanFS file system. In addition to the compute nodes a set of service and administrative nodes provide user login, batch scheduling and management.

Software specification

Operating system: Stokes has Open-SUSE (64-bit) installed as the system software. Hyper-Threading (HT) is enabled on Stokes.

Workload and Resource management: The MOAB scheduler integrated with the Torque resource manager is used for workload and resource management.

Compilers and libraries: Intel C and Fortran compilers are used for compiling the MVAPICH2 and Intel Math Kernel Library (MKL). NPB programs are compiled using the Intel compilers.

Developments

In addition to the environment configuration, job submission scripts (i.e. PBS scripts) are developed for each NPB program run.

Amazon Elastic Compute Cloud (EC2)

The following EC2 instance type is compared against the Stokes compute node. The EC2 cluster computer instances provide proportionally high CPU resources with increased network performance and are well suited for HPC applications and other demanding network-bound applications. The EC2 cluster compute instances are of two types i.e. Quadruple Extra Large instance (i.e. cc1.4xlarge) and Eight Extra Large instance (cc2.8xlarge). We have selected and configured cc1.4xlarge instance type with the following hardware and software specifications.

Hardware specification

Each cc1.4xlarge compute instance is comprised of 2 x Intel Xeon X5570, quad-core "Nehalem" architecture, 23 GB of memory and 1690 GB of storage space. The nodes are interconnected via 10 Gigabit Ethernet. In order to have persistent storage across all running instances Amazon Elastic Block

Storage (EBS) is used. Amazon EBS volumes are off-instance storage that persists independently from the life of an instance. In addition to the compute node, a login node is setup for job submission.

Software specification

- EC2 Instance management: StarCluster is an open source cluster-computing toolkit for the EC2. StarCluster is used to facilitate the process of building, configuring, and managing HPC instances on the Amazon's EC2 cloud. StarCluster provides a set of Amazon Machine Images (AMIs) with a minimal set of software required to setup the HPC environment within the EC2. Each AMI has its own operating system (e.g. Ubuntu, Red Hat Linux etc.) and set of libraries (e.g. OpenMPI).
- Operating System: Ubuntu (64-bit) based AMI is used to setup the instance. Hyper-Threading is enabled by default on the EC2 instances.
- Workload and resource management: The Sun Grid Engine (SGE) is used for scheduling and resource management as it is bundled as the default option for scheduling and resource management in the StarCluster AMIs.
- Compilers and libraries: Evaluation versions of the Intel C and Fortran compilers are used for compiling the MVAPICH2 and Intel Math Kernel Library (MKL). NPB programs are compiled using the Intel compilers.

Developments

In addition to the environment configuration, job submission scripts (i.e. SGE scripts) are developed for each NPB program run.

Table 5.2 summarises the specification for Stokes and EC2 compute nodes:

	Amazon EC2	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, Each compute unit has two Intel (Westmere) Xeon E5650 hex-core processors, 64-bit platform
Connectivity	I/O Performance: Very High (10 Gigabit Ethernet)	ConnectX Infiniband (DDR) providing high bandwidth and low latency for both computational communications and storage access.
OS	Ubuntu (64-bit), HT-enabled	Open-SUSE (64-bit), HT-enabled
Resource manager	Sun Grid Engine	Torque
Compilers & Libraries	Intel C, Intel Fortran, Intel MKL, Intel MVAPICH2	Intel C, Intel Fortran, Intel MKL, Intel MVAPICH2

Table 27 – Stokes and EC2 - System specification

After providing an overview of the benchmarking environment, the next section discusses the first set of benchmarking results.

5.4 NPB Execution and Results

Each NPB program was executed 22 times on both Stokes and EC2 and then results are averaged to generate the graphs (Figure 5). For MPI version of the NPB, the BT and SP programs are compiled for 16 processors and rest of the programs are compiled for 32 processors. The MPI jobs for BT and SP were executed using 2 compute nodes on both Stokes and EC2. For rest of the programs, 3 nodes on Stokes and 4 nodes on EC2 are allocated with 12 and 8 processors per node respectively.

For the OpenMP (OMP) version of the NPB benchmark, 8 OMP threads were specified for all the programs on both Stokes and EC2.

Figure 5 illustrates the MPI and OMP benchmark results for both Stokes and Amazon EC2:

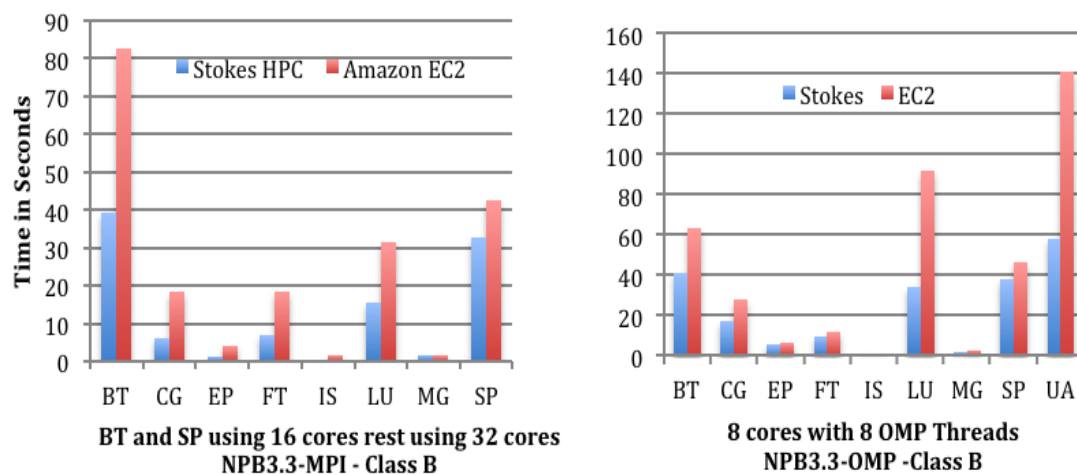


Figure 5: NPB MPI and OMP Benchmark Results³⁷

The preliminary findings, mainly useful to establish a basic range of potential performance multipliers for cost comparisons, 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).
- For the MPI version of the NPB programs, the EC2 performance lags behind the Stokes cluster. The differentiating element here is the 10 Gigabit Ethernet versus the InfiniBand interconnect

³⁷ Please note that although Hyper-Threading was enabled on both the Stokes and EC2, the number of cores refer to actual number of physical cores as opposed to the number of logical cores.

used by the EC2 and Stokes respectively. The average performance loss was 48.42% (ranging from 1.02% to 67.76%).

The above results are for a standard problem size; still the performance differences are obvious in comparing the virtualised general-purpose HPC infrastructure against the purpose-built system for HPC. In addition to performance degradation, the configuration overhead for the EC2 is an additional factor that should be considered.

5.5 Summary and Future Work

In this section, we have covered the strategy to benchmark the HPC, HTC and corresponding offerings from the Cloud computing paradigm. As a first step, we have executed the NPB test case for an HPC system (i.e. Stokes) and the compute cluster instance from the EC2. The results have highlighted that even for an average problem size, the performance degradation is significant. Therefore, additional cost factors in terms of performance penalties and configuration overhead will need to be considered while estimating the cost for various e-infrastructures. The next step is to develop a similar test case for the HTC and corresponding Cloud instances to identify additional cost factors for the HTC and Cloud infrastructures.

6. CONCLUSIONS AND FUTURE PLANS

6.1 CONCLUSIONS

The main conclusions can be summarized as follows:

- Cloud and general commoditisation of computing has acted as a catalyst for cost assessment of e-Infrastructures by providing apparent opportunities for direct comparison. However, as e.g. the cost breakdown analysis and benchmark testing show, the components that can be directly compared are quite limited in number. It is very easy to end up matching fundamentally different resources and services in the same group, with misleading results. However, due to the growing scale and scope of e-Infrastructures, continued cost assessment and comparison efforts will become more and more important.
- Reviewing the literature in comparing clouds with in-house e-Infrastructures shows ratios of maximum 7.22 – 5.59 (the cloud being more expensive) to minimums of around similar prices for reserved instances.
- e-FISCAL initial findings are in line with the above literature, being closer to the lower end of ratios and with the same order of magnitude for €/CPU hour
- Both the cost of dedicated HTC/HPC centres' costs and price of Cloud services are dropping
- The cost is only one of the factors influencing the choice between in-house and outsourced solutions and also one of the aspects affecting the e-Infrastructures sustainability.
- The cost assessment process is painful and laborious, yet essential and useful given all the insights and the learning experience that goes along with it.
- Striking a balance between level of detail and easiness to answer a survey instrument (questionnaire) was challenging, but we had to go for a relatively “ambitious” questionnaire in the first year. We plan to develop a “lighter” version in the second year of the study in order to stimulate increased participation.

Some more detailed conclusions are included for the different sections:

- State-of-the-art review: Although a couple of years ago there were few articles on the topic, there is a lot of recent literature on the costs of individual “in-house” developed HPC or HTC centres or campus Grid systems and their comparison with commercial “on-demand” cloud services, and in particular the dominant player in the market, i.e. Amazon Elastic Compute Cloud (Amazon EC2). The majority of them come up with a price for core/hour, which ranges from around 0.015 Euros (Magellan report) to around 0.075 in the UK (for the JISC study).
- The sample for the e-FISCAL findings was relatively good; 26 answers from 14 countries. However, high-end HPC centres (such as the PRACE Tier-0s) or other high-end HTC centres (such as the WLCG Tier-1s participating in EGI) are not included. This was primarily due to confidentiality reasons and specific non-disclosure agreements between the vendors and those centres that are preventing the publication of detailed cost information.

- The e-FISCAL median values (which take into account outliers) are around 0.05 in 2010 and 0,0337 in 2011, while averages are around 0.1 in 2010 and 0.08 in 2011). This shows that the e-FISCAL initial findings are in-line with the ones reported elsewhere.
- The significant differences between median and average show that there are outliers in the sample (low or high numbers) that significantly influence the averages.
- Costs between 2010 and 2011 are dropping in-line with the trend of lower hardware prices and better overall efficiency. The breakdown between CAPEX and OPEX in 2011 in our calculations is around 30%-70% (median) to 32%-68% (average). Around 50-55% of total cost is dedicated to personnel. The average utilization rate used to calculate the average and median cost per logical/CPU hour for the above results for 2011 is 58% and 67% respectively. This refers to a mixture of EGI, PRACE and other sites not integrated in these e-Infrastructures. As an example, for 2011, EGI reports a utilisation rate of 71,3%. Obviously the higher the utilization, the lower the cost, but in the initial analysis “conservative” values were selected consistently Other interesting findings are the high numbers of depreciation rates for the hardware (average 5 years), the quite good rates of PUE efficiency (of around 1.5 median value) and the percentage of electricity cost (around 15% median value of all costs).
- Comparing with commercial on-demand prices and Amazon EC2 is not straightforward due to several reasons; there is no performance normalization (benchmarking “sanity” efforts not yet concluded), network and storage costs for Amazon need to be taken into account for a complete comparison (note though that network cost information from e-FISCAL questionnaires was also minimal), there will be still some personnel (such as application developers and administrators) that will be involved in the operation of EC2 instances and the adaptation of the application code) that needs to be taken into account, comparing in-house costs with EC2 prices and in fact costs from 2010 and 2011 with EC2 prices in 2012 (that may include a profit or loss) are some of them. Having said that, the pure computing EC2 cost of an “on-demand” instance is 0.06 Euros and for a Linux 3-year reserved instance at 100% utilisation rate it is 0.025 Euros. However, it is improbable to achieve in practice such a utilization rate and the lower the actual utilization rate, the higher the actual cost.
- Benchmark HPC, HTC and commercial Cloud costs (Amazon EC2) is a small-scale effort in the project acting as a “sanity” check. As a first step, the NAS Parallel benchmark has been run in both an HPC system (i.e. Stokes centre in Ireland) and the compute cluster instance from EC2. The results have highlighted that even for an average problem size, the performance degradation is significant and in average around 40%. Therefore, additional cost factors in terms of performance penalties and configuration overhead should also be considered while estimating the cost for various e-infrastructures. As a next step, a similar test case for the HTC and corresponding Cloud instances will be developed to identify additional cost factors for the HTC and Cloud infrastructures.
 - o One issue that was revealed was the not so clear separation between HPC and HTC systems; in fact the low-end HPC systems compared with the high-end HTC system do not have significant differences.

6.2 FUTURE PLANS

This document summarises the initial results of the e-FISCAL project, based on the initial survey results that have allowed the project refine and focus its attention to most critical open questions. More external inputs and validation work is planned for the remaining of the project, namely:

- Discussing these initial results during the e-FISCAL workshop and receiving feedback from financial experts on the interpretation of the results
- Cross-checking the received data. This will be done either during the e-FISCAL workshop in Samos or later on, through interviews or off-line methods.
 - o “Outliers” will be studied and the survey respondents will be asked to provide more information about the background of their values.
 - o Also the issue of physical vs. logical (virtual) cores (i.e. with or without hyper-threading enabled) will be investigated
- Working to get more primary data, enhancing the e-FISCAL sample with more countries and types of centres (such as PRACE Tier-0 ones).
 - o This is being worked out with the coordinators of the EGI-InSPIRE and PRACE IP projects and an effort will be sought to get some minimum information from these centres until the end of the project (through a “lighter” questionnaire). PRACE is also working on the estimation of TCO values (based on the TCO methodology) and we will work on receiving some level of input.
- An attempt to estimate the cost of the EGI and if possible the PRACE infrastructure will be made
- The next step in the benchmarking efforts is to develop a similar test case for comparing the HTC and corresponding Cloud instances. When the benchmarking “sanity” effort will be concluded, a weight factor on the comparison will be applied.
- One or two more workshops are planned, combined with other EGI or PRACE events, towards the end of the project, aiming at both getting missing information and presenting the findings
- More work in the areas of business models and sustainability will be sought, including Green IT best practices
- An investigation to make this effort sustainable, possibly through EGI or PRACE projects will be made. EGI is already working on its first compendium and this seems a good way to make it sustainable. e-IRG is also another possibility that will be explored. As part of this effort, the project will prepare an excel spread sheet (or similar document) with the structured costs that can be used to help centres track costs.

REFERENCES & RESOURCES

Carlyle A. G. et al, (2010), "Cost-effective HPC: The community or the cloud", IEEE Conference on Cloud Computing Technology and Science.

Cohen, S. and Fotis Karagiannis, e-IRGSP2 deliverable D2.4 available at www.e-irg.eu/publications

Cordis (2011), Workshop Report "Towards A European Roadmap for Datacentres and their Economic and Environmental Impact", held in Brussels in November 2011.

EGEE, (2008), "An EGEE comparative study: Grids and Clouds - Evolution or revolution", Members of the EGEE-II Collaboration, 30/05/2008.

EGI-InSPIRE (2012), EU deliverable: D2.30, "Seeking New Horizons: EGI'S role in 2020" Members of EGI-InSPIRE Collaboration, 27/4/2012.

EGI-InSPIRE EU deliverable D2.6 - Integration of clouds and virtualisation into the European production infrastructure, Members of the EGI-InSPIRE Collaboration, 1/3/2011.

EGI-InSPIRE EU deliverable D2.7 - EGI sustainability plan, Members of the EGI-InSPIRE Collaboration, 31/03/2011

Hammond, M., R. Hawtin, L. Gillam and C. Oppenheim (2010), "Cloud Computing for research", Final Report, Curtis + Cartwright Consulting Ltd (7/7/2010).

Hawtin et al. (2012), "Cost analysis of cloud computing for research", Final Report to EPSRC and JISC, 22/2/2012, Curtis+Cartwright Consulting Ltd.

HEP-SPEC Benchmark webpage, <http://hepix.caspar.it/benchmarks/doku.php>.

Hopkinson, L. and P. James, "ICT Related Energy Use, Costs and Carbon Emissions in UK Universities and Colleges - Results from Use of the JISC/SustelIT Footprinting Tool, April 2012, University of Bradford and SustelIT.

Intel-C available at: <http://software.intel.com/en-us/articles/intel-compilers/>.

Intel-Fortran available at: <http://software.intel.com/en-us/articles/intel-compilers/>.

Intel-MKL available at: <http://software.intel.com/en-us/articles/intel-compilers/>.

Koomey et al. (2009), "Assessing trends in the electrical efficiency of computation over time", CA: Analytics Press, Oakland.

Koomey et al. (2011), "Growth in Data centre electricity use 2005 to 2010", CA: Analytics Press, August 1st 2011, Oakland.

Lin A. and Chen (2012), "Cloud computing as an innovation: Perception, attitude, and adoption" International Journal of Information Management, <http://dx.doi.org/10.1016/j.ijinfomgt.2012.04.001>.

Magellan Report on Cloud Computing for Science (2011), U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research (ASCR), December, 2011 (http://science.energy.gov/~media/ascr/pdf/program-documents/docs/Magellan_Final_Report.pdf)

Marston et al. (2011), "Cloud computing — The business perspective" Decision Support Systems, Vol. 51, pp. 176–189.

Misra, S.C and A. Mondal, (2011), "Identification of a company's suitability for the adoption of cloud computing and modelling its corresponding Return on Investment", *Mathematical and Computer Modelling, Volume 53, Issues 3-4*, pp. 504-521.

Moab Scheduler: <http://www.adaptivecomputing.com/products/hpc-products/moab-hpc-basic-edition/>.

MVAPICH2: MPI over InfiniBand, 10GigE/iWARP and RoCE available at: <http://mvapich.cse.ohio-state.edu/overview/mvapich2/>.

NAS Parallel Benchmark Technical Report, available at <http://www.nas.nasa.gov/assets/pdf/techreports/1994/rnr-94-007.pdf>.

Nazir, A, S- A. Sørensen (2010), “Cost benefit analysis of high performance computing infrastructures”, Paper presented in Service Oriented Computing and Applications (SOCA) 2010 IEEE International Conference.

Opitz, A., H. König and S. Szamlewska, (2008), “What Does Grid Computing Cost?” *Journal of Grid Computing*, Vol. 6, No. 4, pp. 385-397.

Rasmussen, N. (2011), “Determining Total Cost of Ownership for Data Center and Network Room Infrastructure, White paper 6, Schneider Electric.

Risch, M and J. Altman, “Cost Analysis of Current Grids and Its Implications for Future Grid Markets”, J. Altmann, D. Neumann, and T. Fahringer (Eds.): GECON 2008, LNCS 5206, pp. 13–27, 2008.

Smith, P., (2011) A Cost-Benefit Analysis of a Campus Computing Grid, *Purdue University*.

StarCluster cluster-computing toolkit available at: <http://web.mit.edu/star/cluster/>

Sun Grid Engine now Open Grid Scheduler available at: <http://sourceforge.net/projects/gridscheduler/>

TORQUE resource manager, Garrick Staples, SC '06: Proceedings of the 2006 ACM/IEEE conference on Supercomputing, ISBN 0-7695-2700-0

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 30/11/2011)

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