


State-of-the art and e-FISCAL survey


Sandra Cohen, AUEB

Matti Heikkurinen, ETL




Not all about cost in the end, impossible to start without

- **Ongoing** sustainability discussion related to e-infrastructures
 - Complex and multifaceted issue, qualitative and quantitative issues
- e-FISCAL is focusing on costs
 - Also performance **adjusted** costs (through benchmarking)
 - Therefore we touch upon only one aspect of sustainability
- Sustainability linked with value produced at least as much as with costs
 - Need both **'R'** and **'I'** to discuss **ROI**



Presentation outline

- e-FISCAL challenges
- Methodology followed
 - State of the art conclusions
- e- FISCAL survey instrument
- Preliminary findings
 - Discussion
- Conclusions
- Next steps



e-FISCAL challenges

- **Getting** cost and infrastructure **data** from HTC and HPC centers
 - Identify similarities, pinpoint differences, model cost patterns, understand cost structures and cost drivers
 - The goal: generic cost model
 - Confidentiality issues emerged
 - Part of the process? Top-500 struggled with similar issues
- Calculate an average cost per logical CPU to allow:
 - Crude comparison with cloud commercial offerings
 - Cross checking the costs in European HTC and HPC centers with published data (mainly US and UK cost studies)

Methodology overview

We are here!

State-of-the-art
review in costing
issues



Collection of data,
Cross-checks
Benchmarking
Conclusions-findings

Development of a
cost model

Sample
identification

Questionnaire
development

Questionnaire
dissemination,
follow up

We have gone through the first full cycle of the methodology and we are about to start again by capitalizing on the feedback and experience gained

This workshop is one of the key steps in the methodology!

State of the art



e-FISCAL project

On going process

39 entries so far

Financial Study for Sustainable Computing
e-Infrastructures

Home | About | Events | Outputs | **State of the Art** | Survey | Contact | FAQ | News

Home

State of the Art

This table provides a list of papers relevant to the e-FISCAL work, which have been reviewed as part of the project state-of-the-art analysis. The papers mainly deal with financial aspects of High Throughput, High Performance, or Cloud Computing or other aspects related to the project such as Energy and Green IT.

If you are aware of a paper that is relevant and not listed below, please send us a note using the [contact form](#). The [e-FISCAL summer workshop](#) (July 3rd - 4th 2012) will also offer a chance to discuss state of the art.

- Multi scope orientation:
 - Costing issues in general
 - Business Models
 - Cloud vs. Grid papers
 - Migration to the Cloud papers
 - Industry benchmarks

News
:: e-FISCAL at the e-IRG workshop
(Copenhagen June 11-12)

- Academic papers,
- Industry project results,
- EU studies' results,
- Engagement with experts
- **No** promotional material

News
:: e-FISCAL computing cost survey-
deadline extended to 29th February

Paper	Summary

Couple of highlights

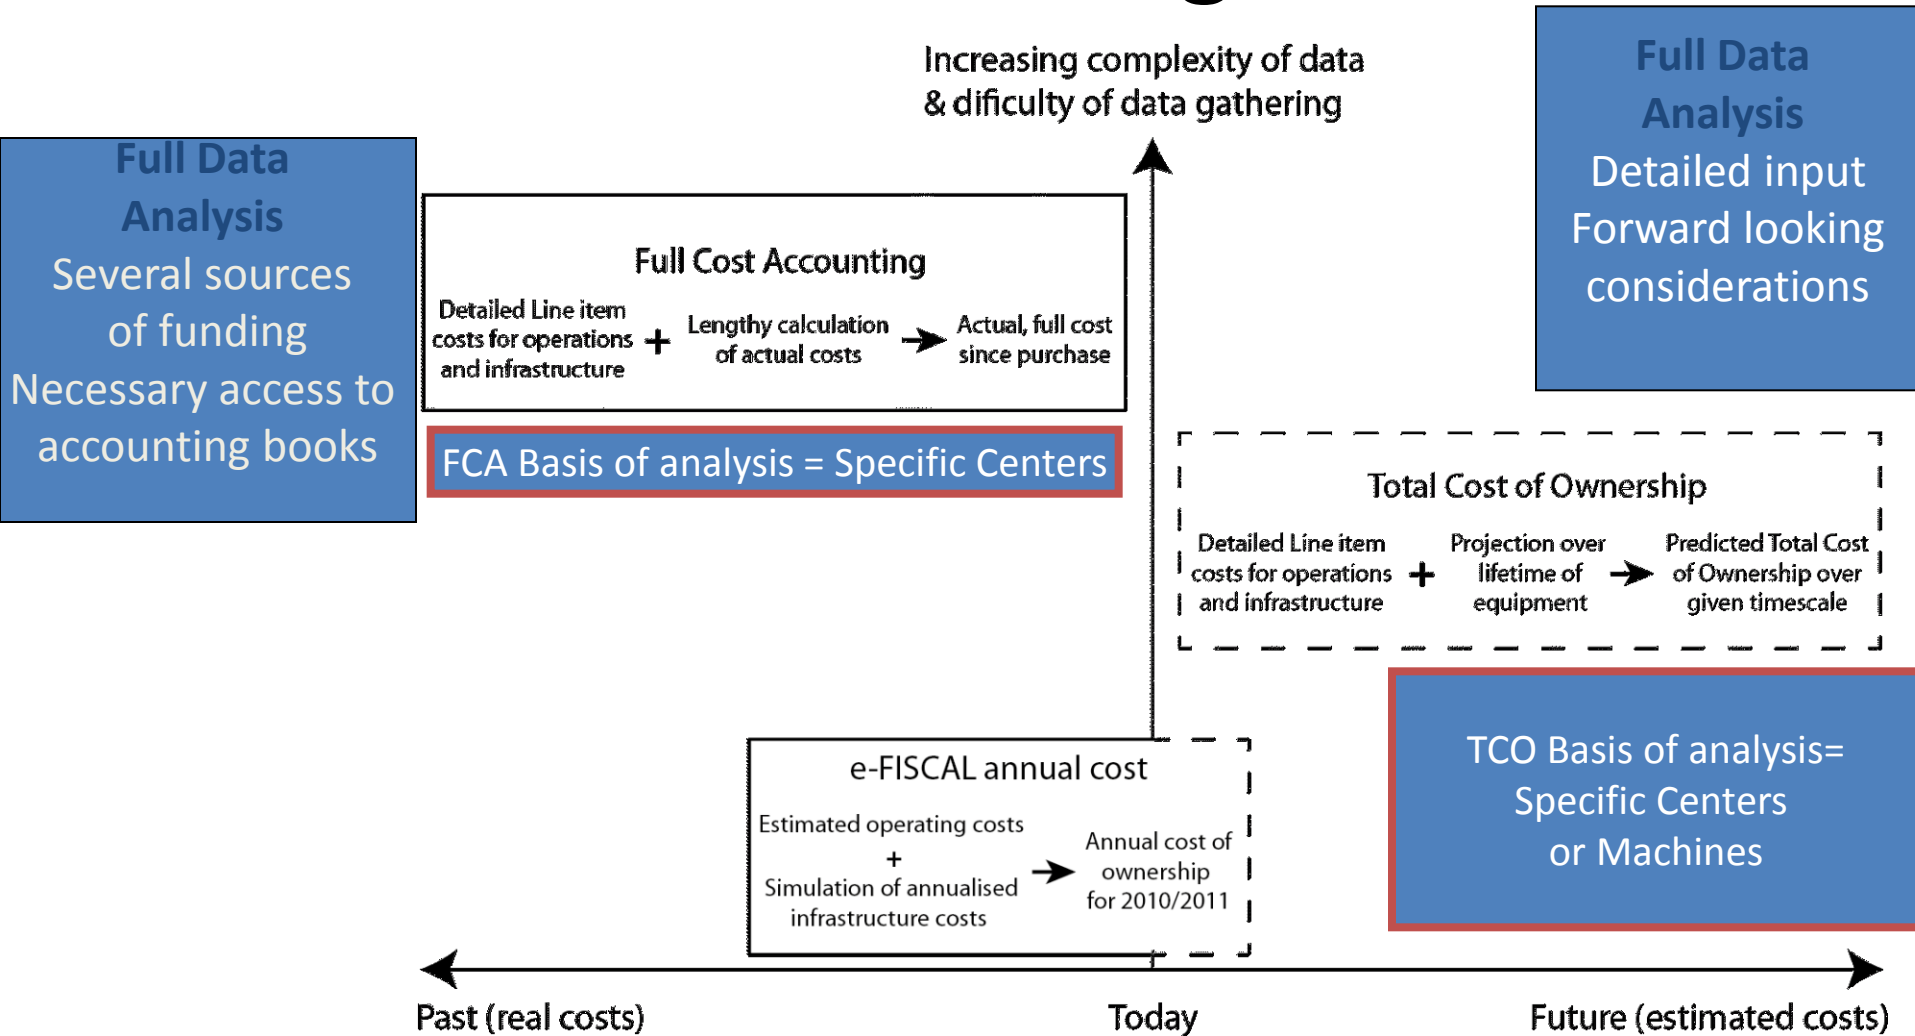
- Hawtin et al. (2012)
 - The more powerful cloud computing instances, rented on an hourly basis, appear to be **one-and-a-half to two times** more expensive per core-hour than well-managed, locally-provided clusters in modern data centres operating at high utilisation levels.
 - However, other purchasing models (such as ‘Reserved Instances’) can reduce the costs **to parity or better**
- Magellan final report (2011)
 - “the cost analysis shows that DOE centers are cost competitive, typically **3-7x less expensive** when compared to commercial cloud providers”

Summary numbers

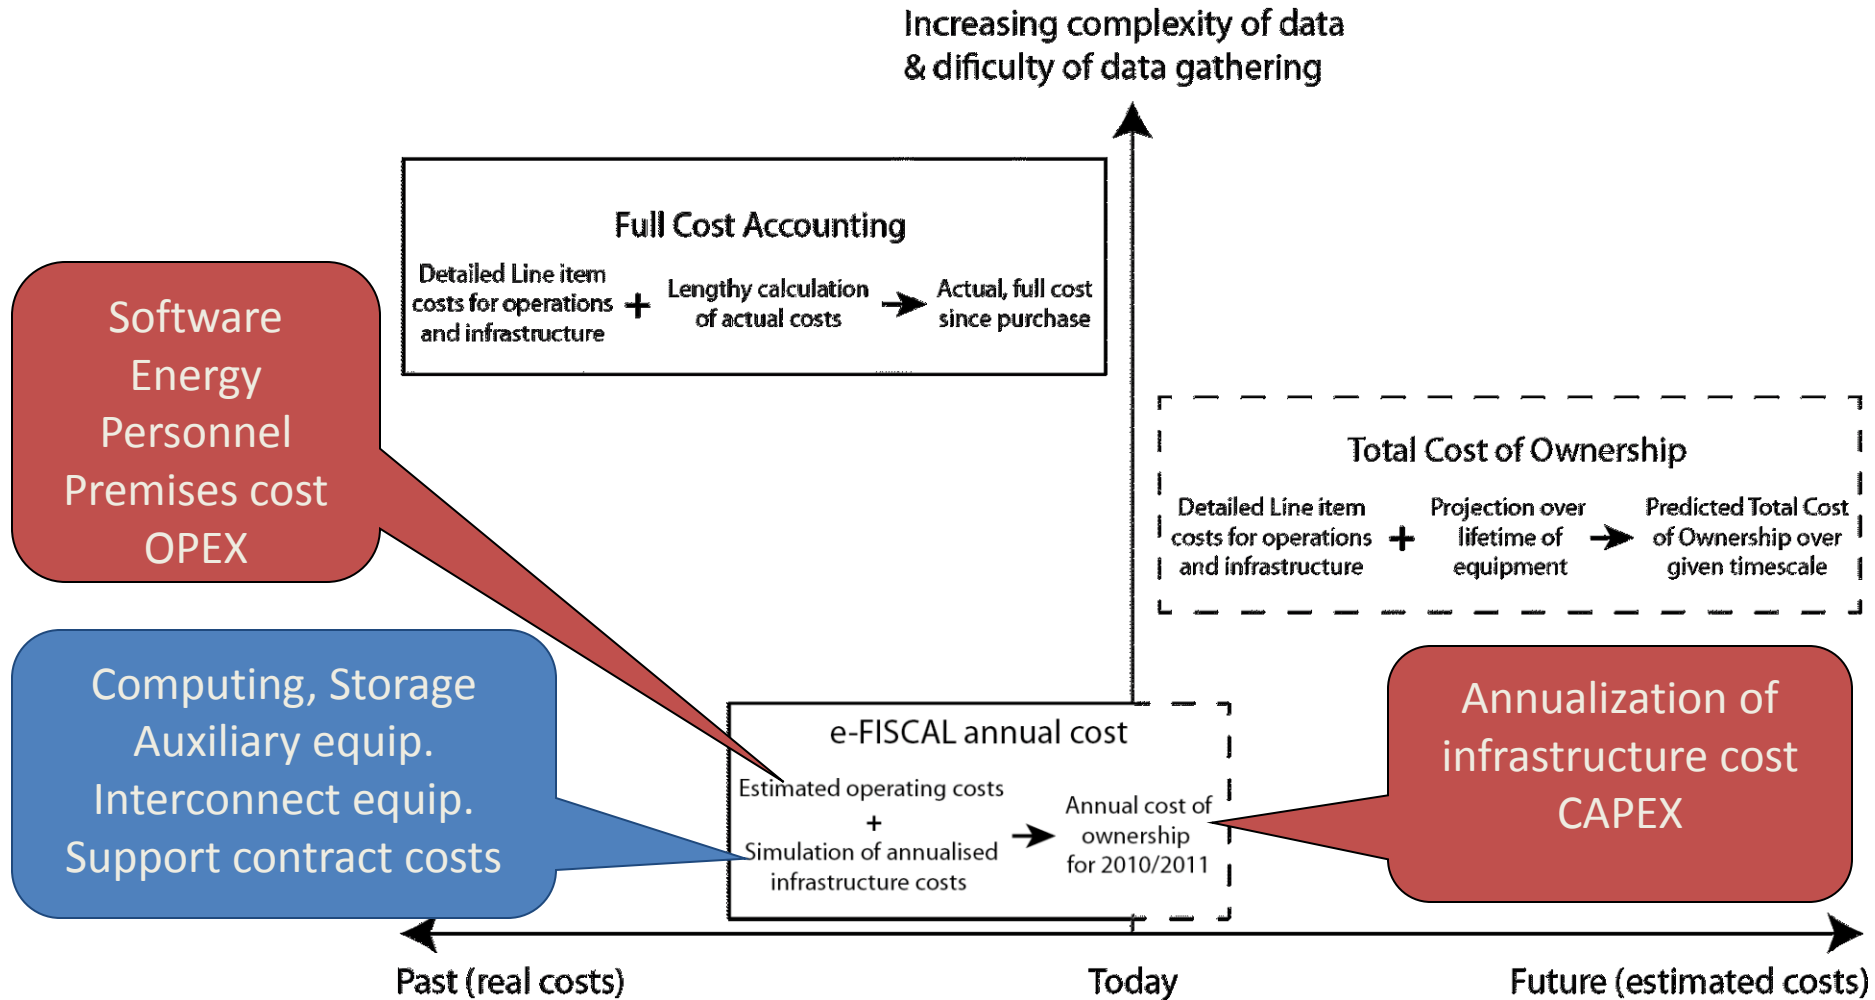
All studies perform a case study or multiple case analysis. e-FISCAL is the first to provide an extended synthesis


Reference	Cost per core hour	Comments
Hawtin et al. (2012)	£0.05 - £0.07	Study for JISC UK - Differences between institutions reviewed
US DoE - Magellan report (2011)	\$ 0.018	Hopper system – National Energy Research Scientific Computing Centre- including storage sub- system
Smith (2011)	\$ 0.039	Purdue campus, USA
University of Washington	\$ 0.051	Hyak cluster, USA
Cohen and Karagiannis (2011)	€ 0.0854 – € 0.1356	Stratified sample of EGI centres - Assuming 60% utilization ratio – storage cost included (numbers refer to 2009)
Cohen and Karagiannis (2011)	€ 0.0782 – € 0.1020	Stratified sample of EGI centres - Assuming 60% utilization ratio – storage cost excluded (numbers refer to 2009)

Basis of costing exercise



Basis of costing exercise



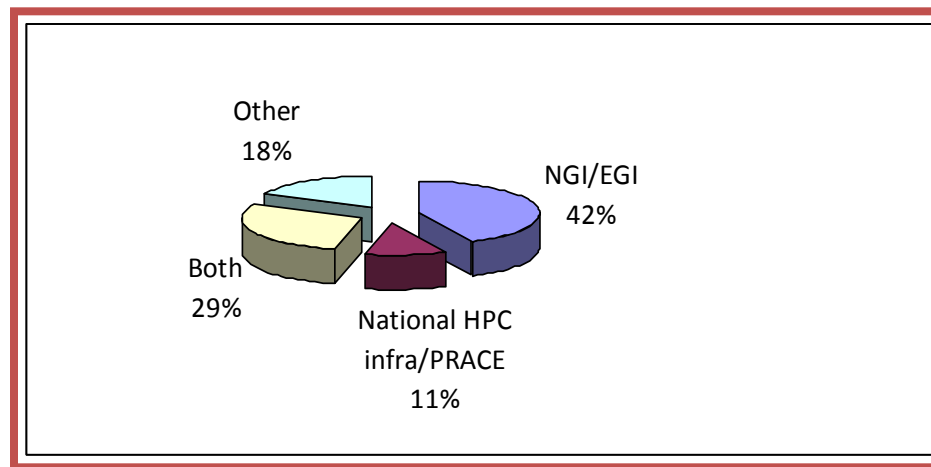
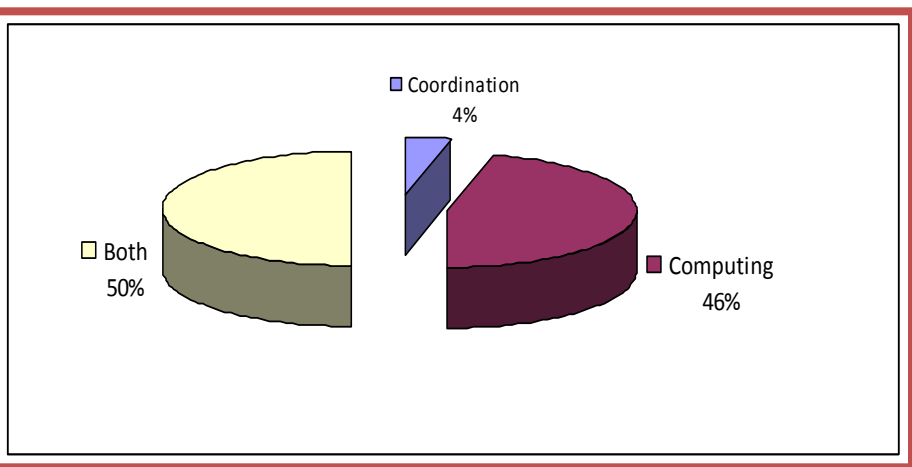


Sample/Respondents so far...

- Ideal sample
 - The total population of:
 - EGI/HTC sites and PRACE/HPC centers
 - Data from several years
- We could do with a fair representation
 - Different HTC/HPC site/centre sizes
 - Countries
 - e-infrastructure types
- Data from 2010 and 2011

Sample/Respondents so far...

- We have gathered information from:
 - 26 respondents – 14 countries




- The vast majority of respondents provide both computing and coordination
- Most of the data from HTC or mixed HTC/HPC centres

Countries contributing

The study will continue



Belgium (5), Bulgaria, Cyprus, Finland, Germany, Greece (4), Hungary, Ireland, Latvia, Norway, Poland, Romania, Spain (6), Turkey



Questionnaire

- 2 versions of the questionnaire
 - On line (surveymonkey)
 - Editable pdf
- Final (official) deadline:
 - End of February 2012
 - Results presented in the Workshop are based on this input
- Couple of general notes
 - Terminology evolving:
 - Issue for HTC/HPC
 - Logical/physical, core/CPU (initial choice - “Logical CPU” - not optimal)

Preliminary findings - Hardware

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.

Answer Options	Min	Max	Average	Median	Answered questions
Cost per logical CPU in € in 2010	100	3150	460	300	17
Cost per TB/ Tapes in € in 2010	50	150	97	94	4
Cost per TB/ Disks in € in 2010	65	6000	704	315	15
Cost per logical CPU in € in 2011	80	3400	439	225	20
Cost per TB/ Tapes in € in 2011	37	125	79	78	4
Cost per TB/ Disks in € in 2011	80	3000	503	250	15

Median mitigates the effect of outliers that influence average metrics

Decreasing trends in costs per logical CPU and Storage per TB

Reluctance to disclose information regarding acquisition costs

Preliminary findings- Useful lives

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

Prolongation of the useful life of computing and storage infrastructure
 Most commonly encountered useful lives in literature for **computing** between 3-4 years
 Depreciation period influences yearly CAPEX.
 The longer the depreciation period the lower the yearly CAPEX
 Less straightforward - obvious effect: Old machines consume more electricity

Preliminary findings – Other infra costs and software

Estimated cost relations of several parameters on computing and hardware storage				
	Min	Max	Average	Median
Related interconnect equipment costs (network devices, cables, etc.) as a percentage of the hardware acquisition cost	0%	30%	10%	10%
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 what would be the investment cost of all auxiliary equipment as percentage of the cost of acquiring computing and hardware storage capacity	5%	35%	17%	20%
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 acquisition cost	0%	15%	4%	2%

Important Cost

Difficult to distinguish from acquisition cost

Very Important Cost difficult to capture


Software enigma CAPEX or OPEX

Preliminary findings - FTEs

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

The salary range is very wide



Plotting 1,000 Logical CPUs and number of FTEs per 1,000 Logical CPUs
Not clear conclusion can be derived

Questionnaires 1 to 19

Preliminary findings- PUE

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
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

Improvement
from 2010 to
2011

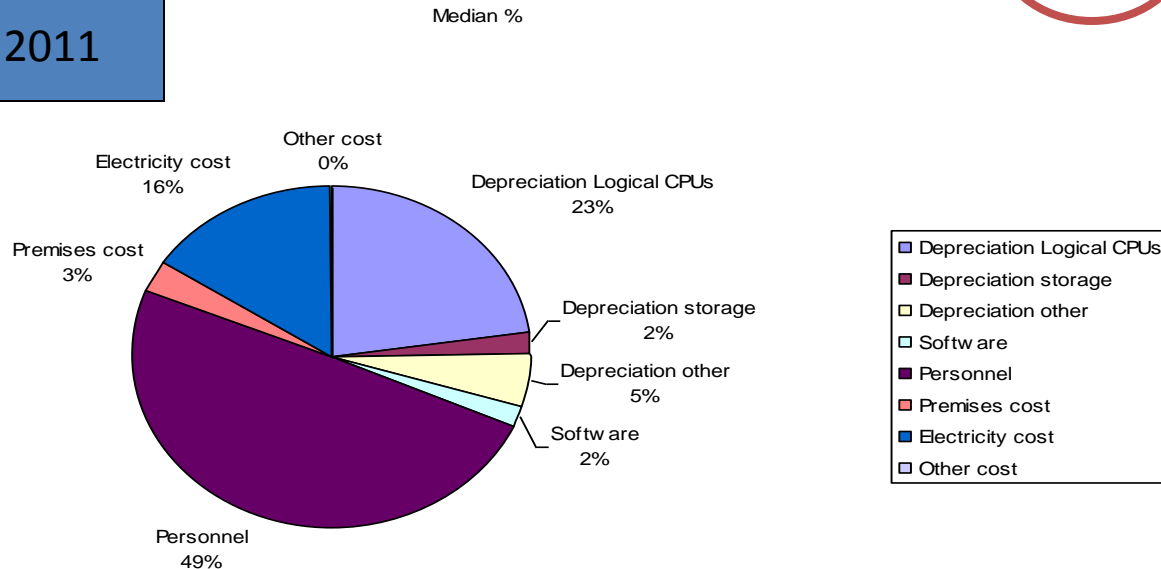
Our respondents were very active in Green IT initiatives (Examples)

- Buying energy efficient servers (improve performance per Watt).
- 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 year.
- Machine rooms in the national infrastructure capture/recycle heat from the compute systems.
- Reallocation of HPC systems.
- Improvement on airflow management
- Implementation of environment monitoring systems

Putting all together

	2010		2011	
	Average	Median	Average	Median
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%

Median 2011



Putting all together

	2010		2011	
	Average	Median	Average	Median
Total yearly cost/ Logical CPU	535.9	258.0	422.7	197.1
Cost per logical CPU/hour	0.1036	0.0499	0.0837	0.0337

Cost per logical CPU/hour is based on a utilization rate calculated from answers to *“logical CPU wall clock time”* and *“number of logical CPUs available at the end of the year”*


This yields conservative numbers, e.g. for median 2011 rate is 67%
Sensitivity analysis based on more realistic numbers and models is ongoing

Utilization rate is a very important factor.

E.g. utilization rate of 80% -> median cost per logical CPU/hour for 2011: € **0.0281**

Depreciation rate is another very important factor.


Depreciation rate 3 years -> median cost per logical CPU/hour for 2011: € **0.0588**



A high level comparison

e-FISCAL vs Amazon reserved instances


Comparison with reserved instances as less costly solutions compared to “on – demand” instances (all amounts in €)



*Cost for 3-year reserved instances/hour transformed in €/logical CPU hour (equivalence based on instance characteristics) Based on windows/EU-Ireland/80% -100% usage of reserved instances. Amazon site accessed on 22/5/2012


**Cost per logical CPU/hour €/logical CPU hour (2011)

No performance adjustment has been performed YET



A high level comparison e-FISCAL vs Amazon on demand

Comparison with on demand instances



*Cost for on demand instances/hour transformed in €/logical CPU hour (equivalence based on instance characteristics) Based on windows/EU-Ireland/heavily used reserved instances Amazon site accessed on 22/5/2012

**Cost per logical CPU/hour €/logical CPU hour (2011)

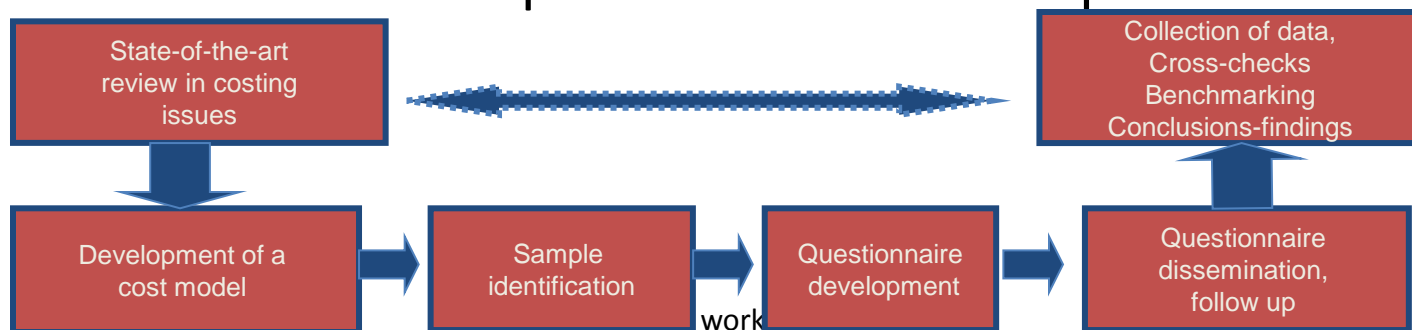
No performance adjustment has been performed YET


Conclusions

- **e-FISCAL novelty:** Assessing and comparing costs under different setting
 - Coming up with an average (median) cost that comprises the diversity
 - Emphasis should be better placed on the identified range of costs
- Our results are inline with literature
 - Cost per logical CPU/hour € 0.0337 (median 2011)
 - Costs show decreasing trends
 - Not only hardware costs but Opex (evidence of existence of economies of scale)
- Nevertheless some interesting issues emerged:
 - Divergence in cost structures
 - Useful lives
 - FTEs/logical CPU and personnel costs
 - Non- unanimous economies of scale existence

Next steps

- Resolving ambiguities in data
- Study methodologies used by sites to come up with energy efficiency ratios and utilization
- Increasing the sample with more respondents
 - Condensed version of the questionnaire
 - Stronger anonymity guarantees
- Combining benchmarking outcomes with cost information
 - Calculation of performance adjusted cost metrics for better comparison with cloud commercial offering
- Collect feedback to improve our model and procedures!





Thank you!

Windows

Amazon EC2 (<http://aws.amazon.com/ec2/pricing/>)

EU (Ireland)	3 yr Term		100% utilization	80% utilization
Standard Reserved Instances	Upfront	Hourly	total/hour (€/Hour)	total/hour (€/Hour)
Small (Default)	300	0,04	0,040	0,043
Medium	600	0,08	0,081	0,085
Large	1200	0,16	0,161	0,170
Extra Large	2400	0,32	0,322	0,340
Micro Reserved Instances				
Micro	100	0,014	0,014	0,015
High-Memory Reserved Instances				
Extra Large	1550	0,198	0,201	0,213
Double Extra Large	3100	0,396	0,403	0,426
Quadruple Extra Large	6200	0,792	0,805	0,851
High-CPU Reserved Instances				
Medium	775	0,12	0,117	0,123
Extra Large	3100	0,48	0,468	0,491

0,081 – 0,085
€/per core

0,100– 0,106
€/per core

Exchange rate \$/€
 0,783208
 0,78320802 <http://www.alpha.gr/tools/eeurocur.htm>
 Accessed 22/5/2012