


Assessing the Costs of the European e-Infrastructure: Mission Impossible?


Matti Heikkurinen
**e-FISCAL Project Consortium/
Emergence Tech Ltd**

e-IRG workshop June 11 – 12 Copenhagen
Danish EU presidency




Numbers: HTC, HPC and Cloud

- Cloud, commodisation catalysts for cost assessment
 - Inevitable with growing scale – also without catalysts
- Cloud/dedicated cost ratios (literature)
 - Maximums: 7.22 – 5.59, Minimums: > 1.00
 - High utilisation rate!
 - Cloud savings hinted, not shown
- e-Fiscal (very) initial results in line with literature
 - Closer to the lower end of ratios
 - Same order of magnitude for €/CPU hour




Better numbers: how hard can it really be?

- Measuring costs is easy
 - Track spending
 - Budget for the known future expenditures
- Comparison is easy
 - Death of the distance, death of the location..
 - Everything is virtualised, CPU hour is a CPU hour
- No major surprises tomorrow
 - Mature technologies
 - Mature business models




Everyday example of the challenge

- Evaluate “LaaS offering” (Life as a Service) vs. traditional, personalised “In-housing” living
- LaaS benefits: flexibility, support labor mobility, optimised commuting and zoning,...
- Money has been used for housing for hundreds of years, easy to compare



Tracking “in-housing” cost

- Let’s ask a sample group their weekly housing costs
 - #1: “who wants to know?”
 - #2: tens of funding sources/person
 - #3: how to deal with gifts, durable consumer goods, rent paid by employer,...
- Need historical data
 - “Grandparents paid 10£ for this table in 1920”
 - Corresponds to £313 - £992 in 2010 pounds
 - Obsolete cost items
 - I paid 500 € (inflation adjusted) to fix my VHS
- Basic challenge with full cost accounting!



Budget for the future


- Predicting is difficult – especially of the future
 - Estimating the (remaining) useful life?
 - When will my fridge break down?
 - Energy costs?
 - Increasing demand, dwindling supply
 - Legal framework
 - Changes in inspection/maintenance regime
 - Building code (asbestos)
- Easy to both over- and underestimate
- Basic challenges with Total Cost of Ownership (TCO)

House is a house?




Be it housing or ICT

- Tracking costs is not easy
 - Past or future
- Comparison is not trivial
 - HPC, HTC, HPC Cloud, HTC Cloud (“house is not a house”)
- Future is uncertain
 - Plummeting prices as a rule, but
 - Flood in Thailand -> HD prices double
 - New technologies
 - Energy costs, green regulations

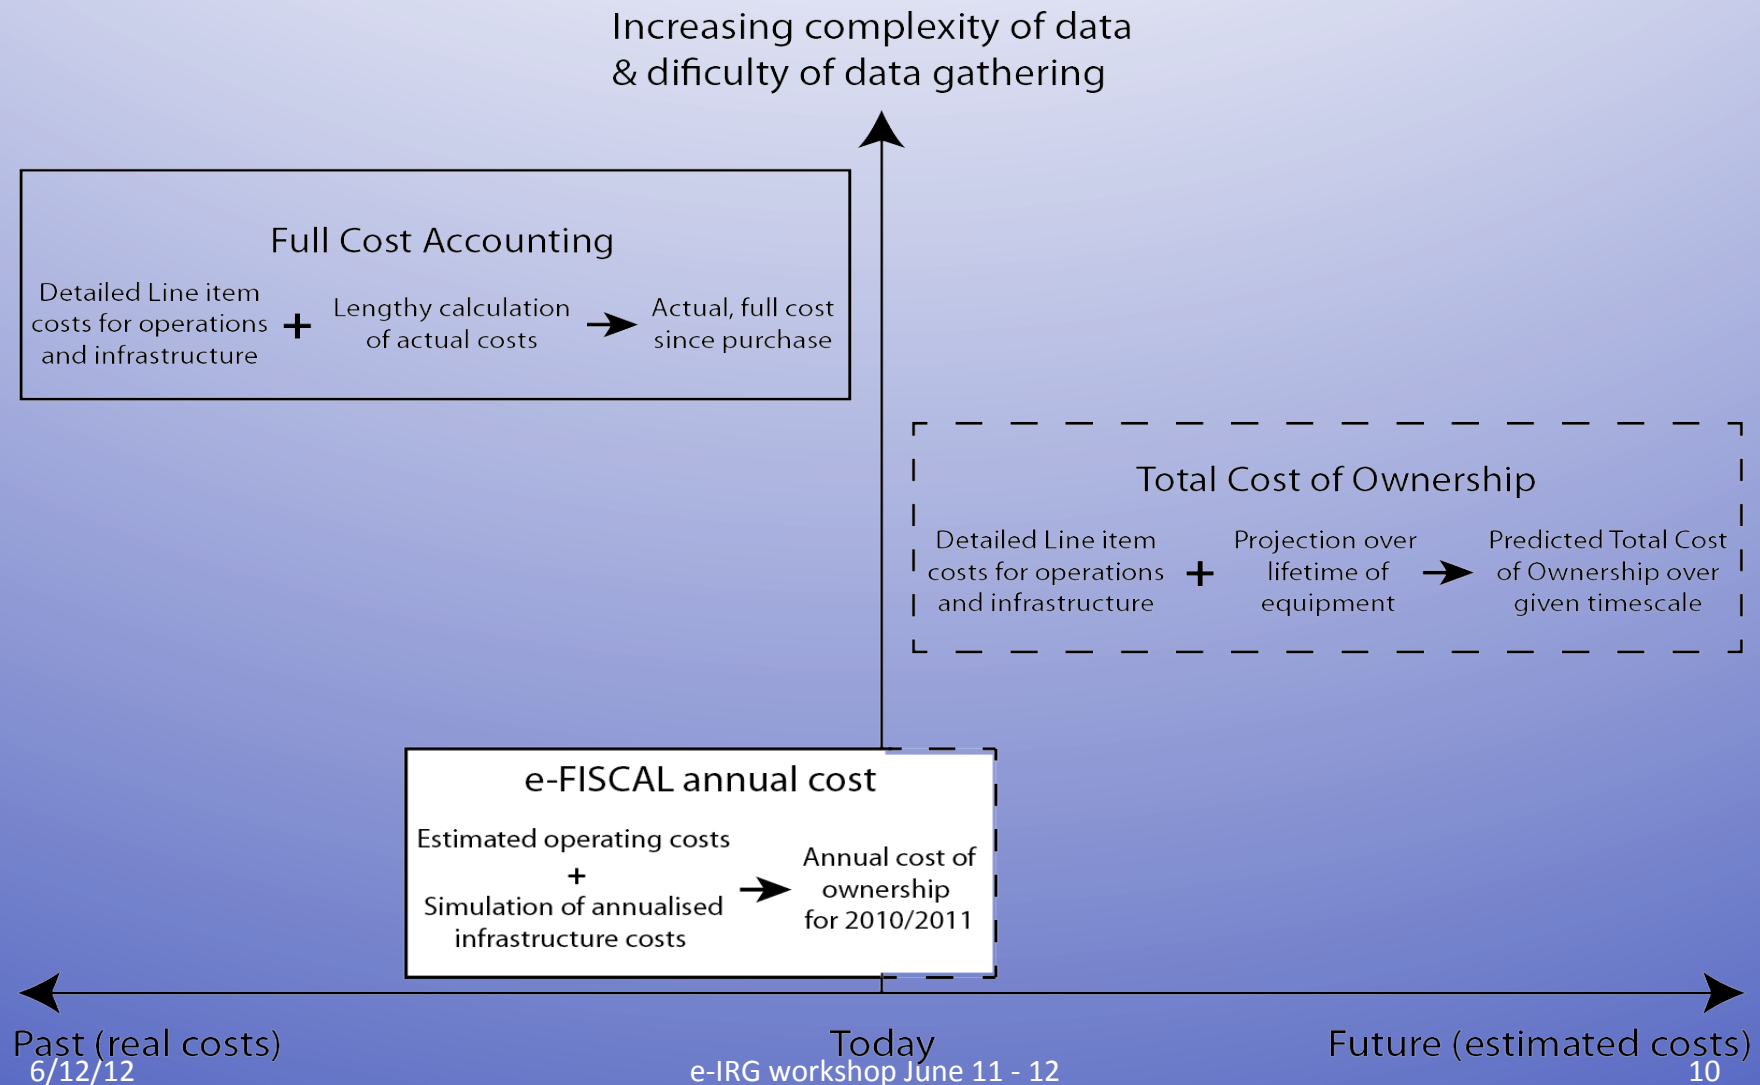


The e-FISCAL approach

- “See the map, determine building costs today”
 - Still need survey data, but considerably less
 - CAPEX can *often* be derived, OPEX needs survey data
 - With ICT we can focus on utility value
 - Old buildings become historic
 - Old computers become recycling challenges
 - We don’t need access to everyone’s bank accounts, nor a crystal ball.






TCO vs. FCO vs. e-FISCAL



Moore's law and others

- *Some* technologies change
 - CPUs, storage systems,...
- Rapid changes in use
 - Capacity increase
 - New usage patterns
- But laws of physics still apply
 - Energy
 - Buildings
 - Speed of light





e-FISCAL activities

1. State of the art survey
 - Ongoing effort
 - Public repository to engage with experts
2. Benchmarking
 - Small-scale, “sanity check”
 - Uncovered HTC/HPC definitions issue!
3. Survey design and execution
 - Deciding on the level of detail a challenge
 - Both financial and technical!

1 - State of the art

- Magellan final report , US DoE (2011)
 - *DOE centers typically 3-7x less expensive compared to commercial Cloud.*
- Carlyle et al. , Purdue University (2010)
 - *Purdue HPC “community cluster” program*
 - *Majority of community: substantially lower out-of-pocket costs per CPU hour*
 - *High utilisation rate -> lower cost in flat fee environment*



1- State of the art

- Hawtin et al. (2012), Curtis+Cartwright for EPSRC and JISC
 - *Cloud on hourly basis: 1,5 – 2 times the cost per core-hour*
 - *'Reserved Instances' can reduce the costs to parity or better*
- Marston et al. University of Florida (2011): qualitative analysis of Cloud
 - *Strengths, e.g. reduced infrastructure costs and energy savings as well reduced upgrades and maintenance costs*
 - *Weaknesses, e.g. the loss of physical control of the data that is put on the cloud*
 - *Opportunities vs. threats*

2- Benchmarking

- NAS Parallel Benchmark – Class B (OpenMP and MPI)

Stokes HPC System	EC2 Cluster Compute Instance
24 GB of memory 2 x Intel Xeon E5650, hex-core processors 143TB of storage, 64-bit platform ConnectX Infiniband (DDR) interconnect	23 GB of memory 2 x Intel Xeon X5570, quad-core processors 1690 GB of storage, 64-bit platform 10 Gigabit Ethernet interconnect



- Average performance loss 37 – 48% (1 to 68% per test)
- Caveat: heterogeneous environments aren't trivial!

3 - Survey sample: 26 answers



3 - Initial cost assessment

	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
Cost per logical CPU/hour €	0.1036	0.0499	0.0837	0.0337

- These shouldn't be compared to Cloud costs today
 - **Outliers**, no performance normalisation, “December resources”,...
 - Both the costs and Cloud prices change rapidly – 2010 costs vs. 2012 prices?
- Having said that: EC2 standard instance hourly rate at today's rates:
 - 0.06€ (on-demand), 0.025€ (Linux 3-year reserved instance at 100% utilisation rate)
 - Need to add storage & extra network costs – often not powerful enough

3 - Sample used

	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

3 - Cost breakdown 2010

		2010		2010	
Cost break down / logical CPU		Average	Median	Average %	Median %
Depreciation Logical CPUs	CAPEX	92,9	60,0	17%	23%
Depreciation storage		26,2	4,2	5%	2%
Depreciation other		29,9	13,5	6%	5%
Software	OPEX	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%

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

- CAPEX fairly small, OPEX dominates (70%)
- Electricity: PUE rates good (median 1.51)
- Depreciation rate important for CAPEX (5 years used based on survey)

3 - Cost breakdown 2011


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

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

- Multi-core technology has a visible impact
- Cloud approach would still most likely need personnel
 - Carlyle: Cloud does not lessen the need for system administration tasks


Conclusions

- Both cost of dedicated HTC/HPC and price of Cloud services dropping
 - “Not quite” Moore’s law
- Cost only one of the factor influencing the choice
 - Convenience, flexibility
 - Intangible factors
 - Non-standard requirements
 - Risk management
- Cost assessment process
 - Painful and laborious
 - Essential
 - Useful (insights, learning experience)



Cost assessment challenges

- Fundamental cost uncertainty principle
 - Partially addressed by e-FISCAL methodology
- Sufficient data rapidly, efficiently
 - Issue for the service provider community
 - Be sensitive to changes “sufficient”
 - Find a trusted party?
- Use the data for policy formation
 - Gain and maintain trust of users, funding agencies
 - Relevance: link cost to value (CPU hours/Nobel price)
 - Forums like e-IRG in an important role



e-FISCAL Summer Workshop

■ 3-4 July 2012, Samos Island, Greece


<http://www.efiscal.eu/2nd-workshop>

@ Samos Summit 2012 event series

<http://samos-summit.blogspot.com>


- 70 registrants (around 30 already expressed interest for e-FISCAL)
- Case studies on Costs-Green IT from:
 - Belgium, Greece, Spain, Poland
 - And LIFEWATCH ESFRI project
- Talks from Intel (tbc) & 451Group
- Benchmarking efforts (HPC vs. Amazon)
- EGI and PRACE talks





Thank you!

Questions, comments?



Project in a nutshell

- **Project acronym:** e-FISCAL
- **Contract n°:** RI-283449 (CSA-SA)
- **Start date:** 01/08/2011
- **Duration:** 18 months (end 31/1/2013)
- **Total budget:** 392.523 €
- **Total funded effort in PMs:** 33.75
- **Partners:**
 - AUEB-RC, EGI.eu, NUI Galway, ETL
- **Web site:** www.efiscal.eu