e-FISCAL

e-infrastructure CAPACITIES

e-FISCAL Final Workshop
Amsterdam, 28 January 2013

Introduction and key findings

Fotis Karagiannis, Sandra Cohen,

Athens University of Economics and Business-Research Center (AUEB-RC)

Evolving EGI Workshop and co-located e-FISCAL Workshop

e-FISCAL

It's all about knowing the costs..

…their composition..

…and putting them in context!

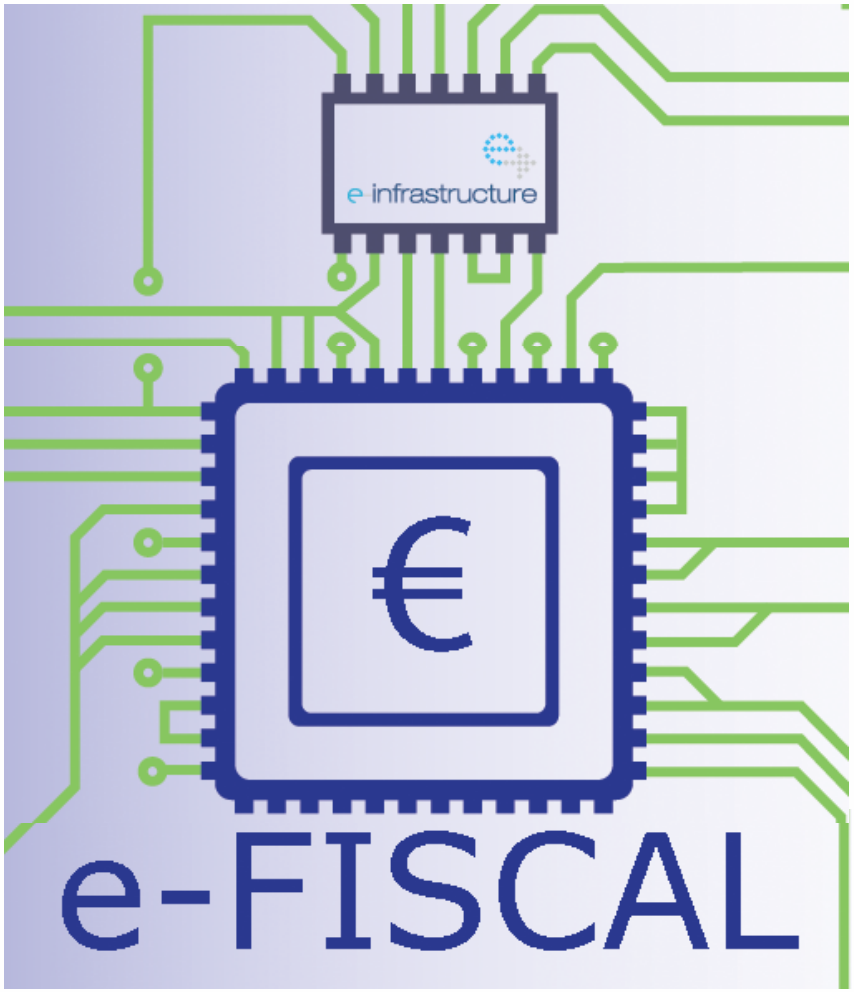

Consortium

AUEB-RC

EGI.eu

NUIG*
(ICHEC)

ETL

* National University of Ireland, Galway / Irish Centre for High End Computing (ICHEC)

Main objectives

- ***Analyse*** *the costs of the current European dedicated High Throughput and High Performance Computing (HTC/HPC) e-Infrastructures for research*
- ***Compare*** *them with the closest equivalent commercial leased or on-demand offerings*
  - *Cloud computing!*
- ***Evaluate*** *the findings through a report*

Background

- First in-depth study at European scale
  - Significant sample of participants, HTC/HPC, comparisons with Clouds, innovative methodology!
- Builds on previous financial exercise
  - e-IRGSP2 project, focusing on HTC/NGIs only
  - More at http://www.efiscal.eu/state-of-the-art

Basis of costing exercise

Full Data Analysis
Several sources of funding
Necessary access to accounting books

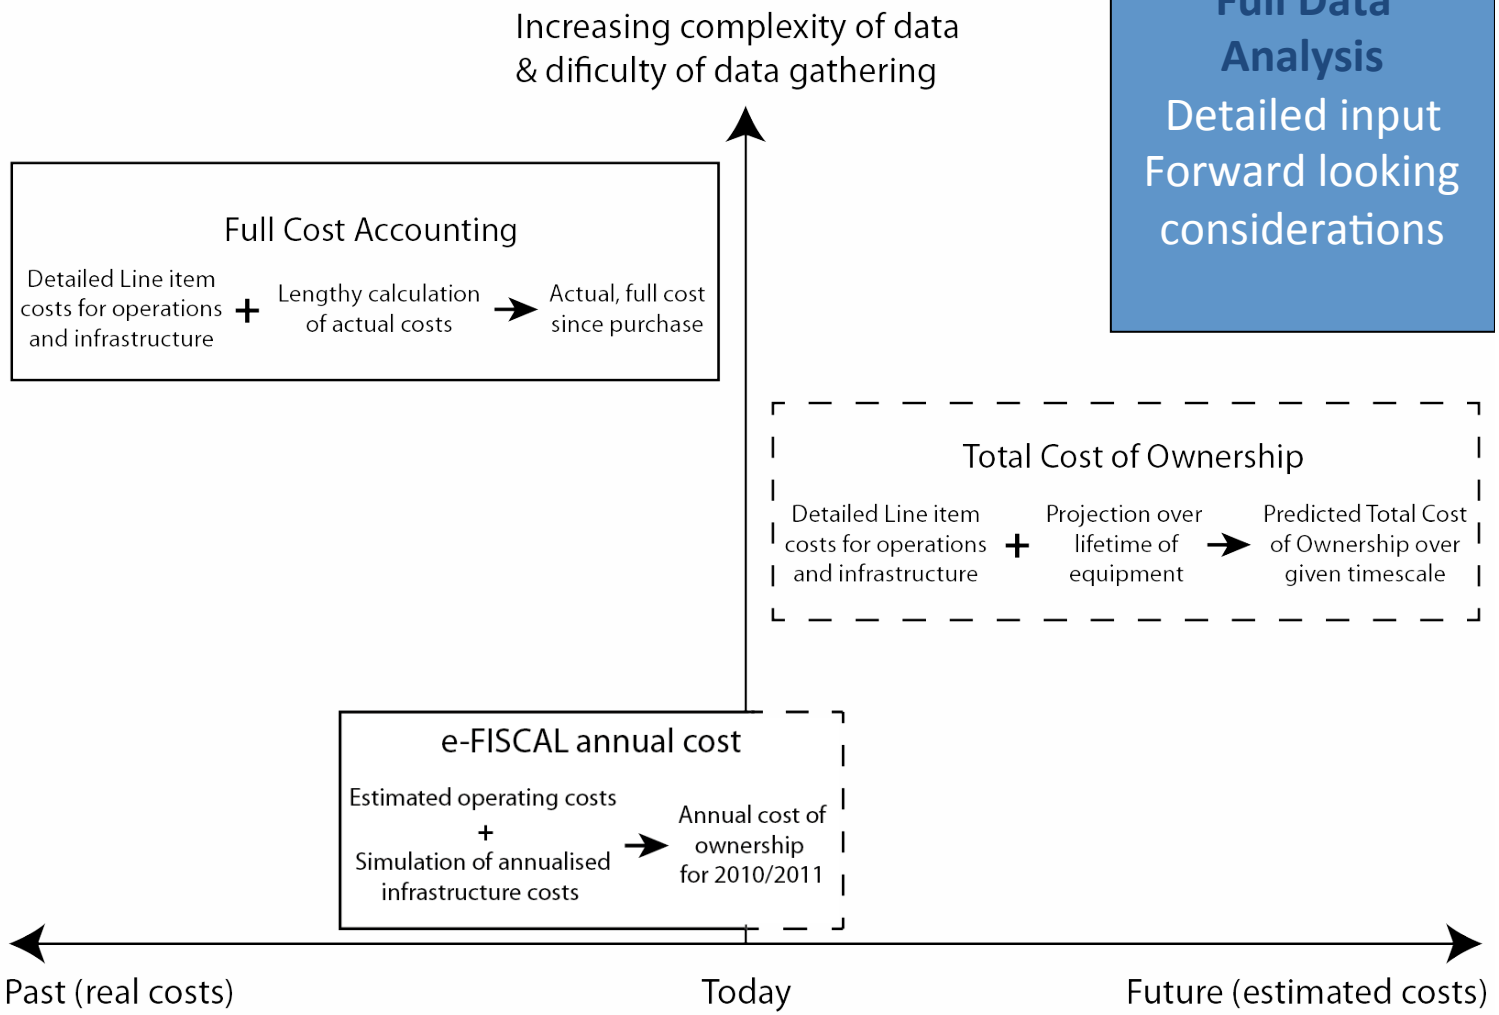

Full Data Analysis
Detailed input
Forward looking considerations

Methodology overview

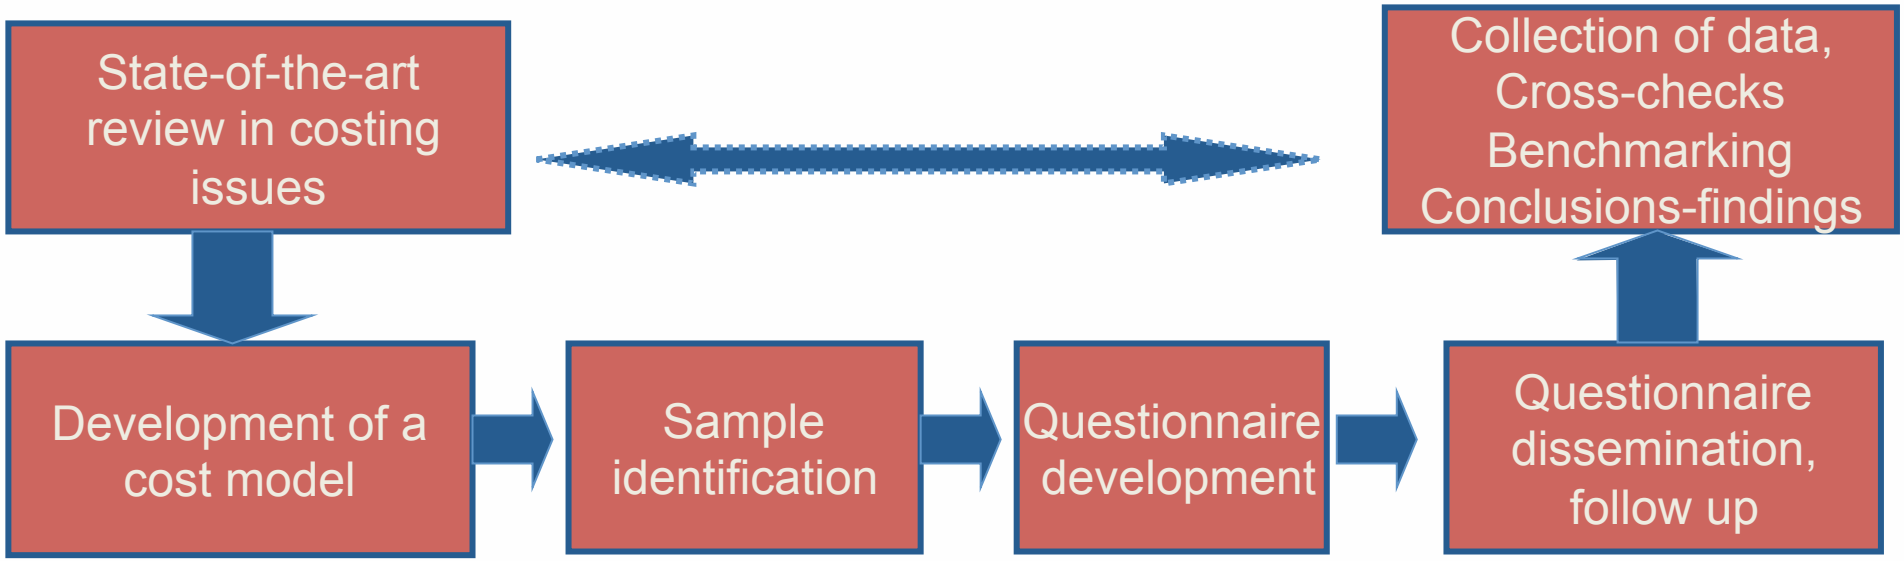

We have gone through the first full cycle of the methodology and we are about to start again by capitalizing on the feedback and experience gained

Methodology

Logical CPUs, storage devices, auxiliary equipment, connectivity devices
X
Prices per logical CPU, for storage, etc.
retrieved by questionnaires

→ **Approximation of the physical infrastructure investment cost**

Development of the financial model

Yearly Operational Expenditures (OPEX)
Software, Personnel, Electricity, Premises, Network connectivity, Other operating costs
(questionnaire)

Yearly Capital Expenditures (CAPEX)
Depreciation of the physical infrastructure costs
(questionnaire)

Total yearly cost of ownership

Not trivial….

- Careful in estimating e-Infrastructure costs and comparing with Cloud prices!
  - Cross-checks/validation with market or other prices
  - Benchmarking comparisons to optimise results
  - Profit-margin possible
  - Moving to the cloud a different exercise!
- Confidentiality/Anonymity of data!
  - No identifiable data related to an individual site or national HPC/HTC entity are presented
  - Some “big” sites (mainly PRACE Tier-0s) not ready to provide data
    - Still PRACE Tier-1s and other EGI big sites participating
- Cost is different from value!

Countries contributing

We would like to thank all contributors!

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

Sample/Respondents

- We have gathered information from:
  - 28 respondents – 16 countries

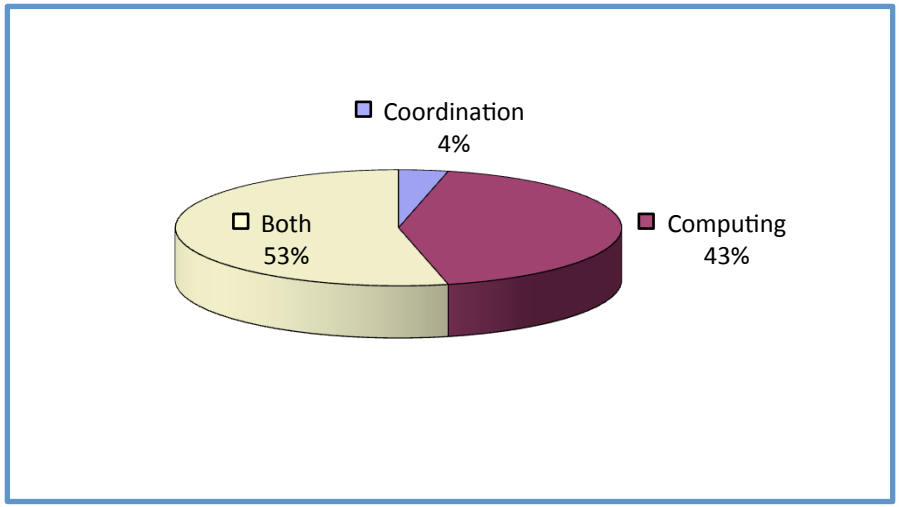

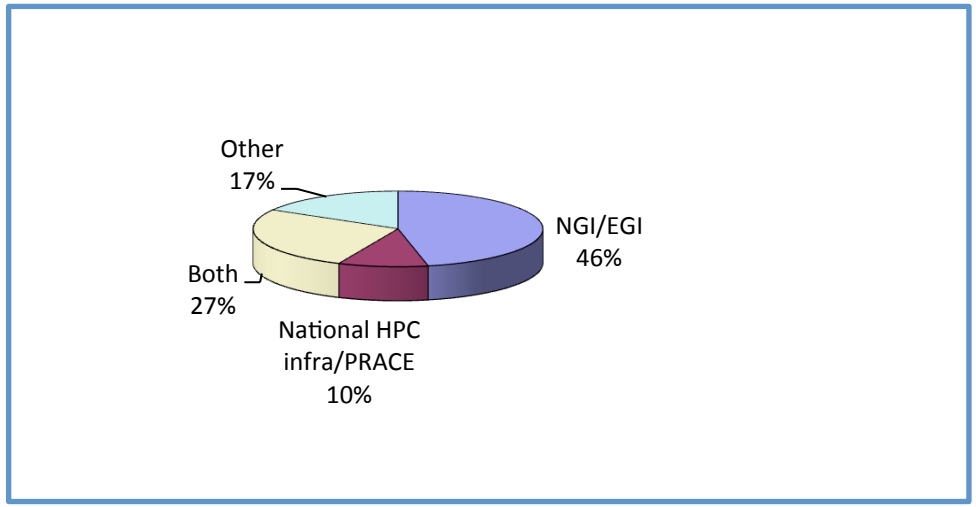

  - The vast majority of respondents provide both computing and coordination
  - Most of the data from HTC or mixed HTC/HPC centres

Review the state-of-th

http://www.efiscal.eu/state-of-the-art

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 (~€0,06-0,09)	Study for JISC UK
US DoE - Magellan report (2011)	$ 0.018 (~€0,014)	Hopper system – National Energy Research Scientific Computing Centre- including storage sub- system
Smith (2011)	$ 0.039 (~€0,03)	Purdue campus, USA
University of Washington	$ 0.025 (~€0,02)	Hyak cluster, USA
Cohen and Karagiannis (2011)	€ 0.0782 – € 0.1020	e-IRGSP2 study: Stratified sample of EGI centres - Assuming 60% utilization ratio – storage cost excluded (numbers refer to 2009)

e-FISCAL main findings

- In-house HPC/HTC e-Infrastructures are cost-effective
  - With high utilisation rates & depreciation rates (as reported)
  - However per application cost analysis is needed
- Personnel ~50% of total costs; CAPEX/OPEX=30/70%
- Larger sites have in general less FTEs/core
- Small-scale benchmarking efforts between in-house HPC and Amazon Compute Cluster instance:
  - A ~40% performance degradation of the latter for HPC, a bit better for HTC (more on next presentation!)
- Modest size HPC centres similar to state-of-the-art HTC ones

More details (1)

	Average	-	Median
• CAPEX / OPEX ratio in 2011:	27/73%	-	**31/69%**
• Personnel / Total costs in 2011:			**50% !**
• Cost per core hour in € in 2011:	0,072	-	**0,031**

Median for minimum utilisation rate: 75%

Likely underestimated, at 80% rate, the cost drops to : €0,029

Depreciation rate: 5 years

For a value of 3 years it goes up to € 0,037

More details (2)

	Average	-	Median
• Cost per core in € in 2011:	390		**204**
• CPU useful lives:	5		**5**
• Interconnect equipment:	10%		**10% of CPUs hw costs**
• Software costs:	4%		**2% of CPUs hw costs**
• Average salary in € in 2011:	53k		**49k**
• Power Usage Effectiveness:	1,55		**1,49**

Costs breakdown
(2011-median)

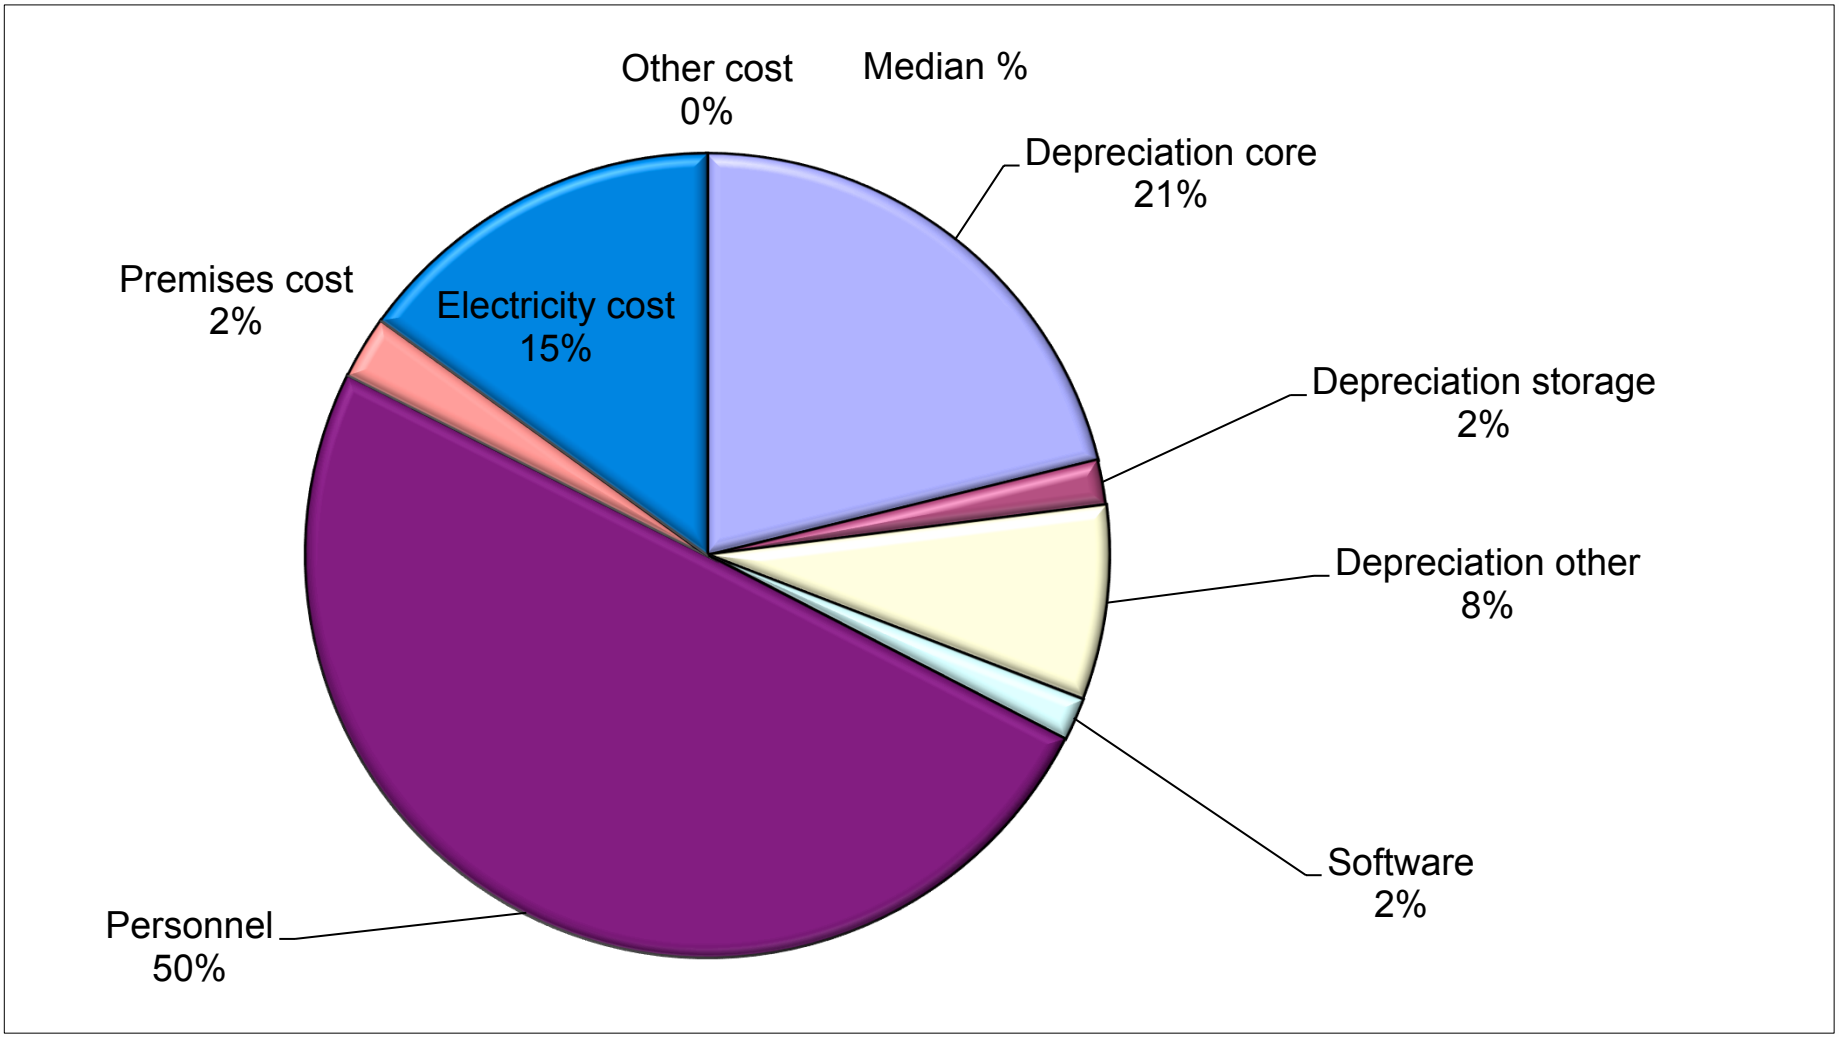

Cost per core hour in € / no of cores*

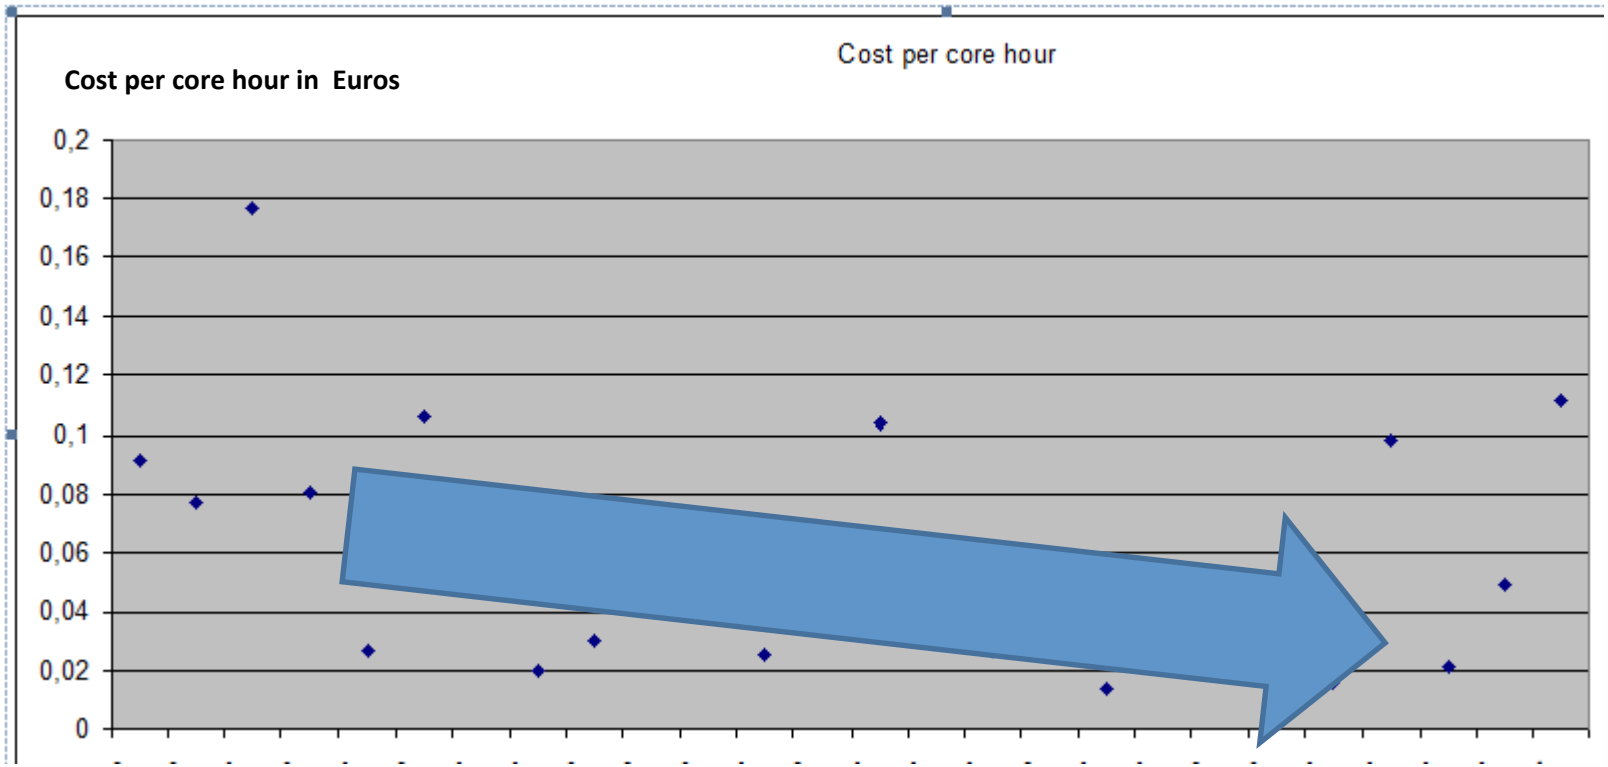

* Dots are sites

Larger sites are in general more cost effective – however outliers exist

NO PERFORMANCE ADJUSTMENTS

e-FISCAL vs. Amazon EC2

e-FISCAL results compared with **EC2 reserved instances** (**all amounts in €**)
Costs refer to 2011 – Prices refer to 1/2013

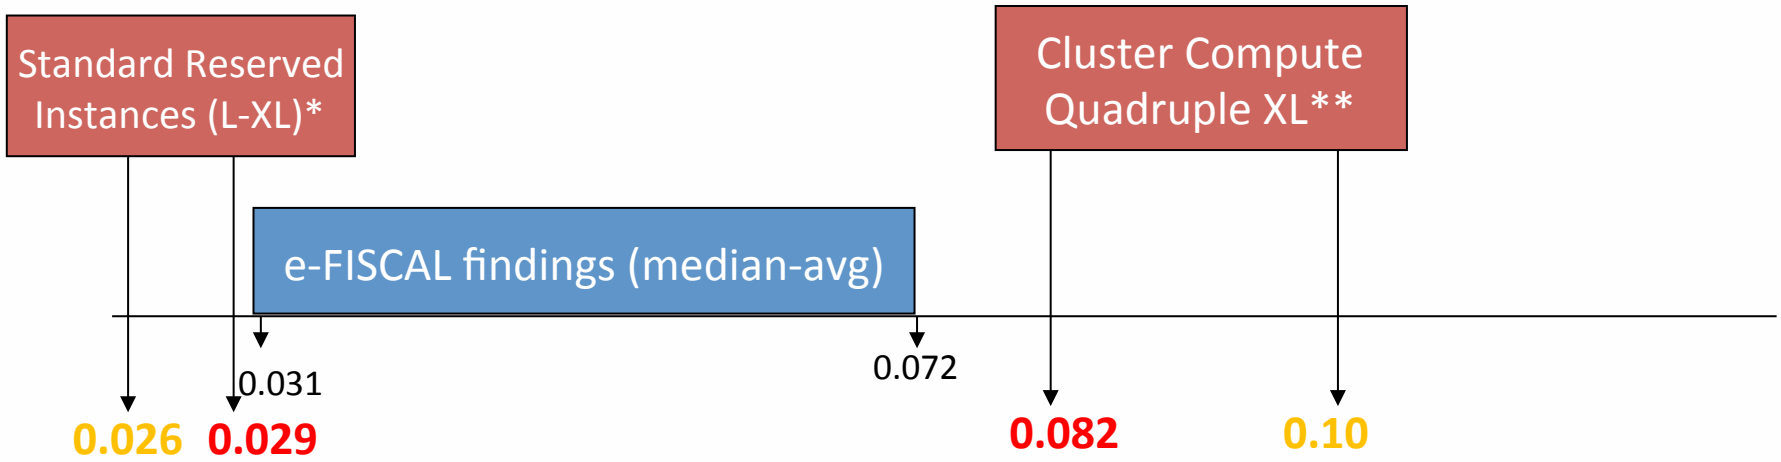

*Price for 3-year reserved instances/hour transformed in €/core hour
(equivalence based on instance characteristics)
Based on Linux 60% (**red**) -80% (**yellow**) usage of reserved instances.
Amazon site accessed on 15/1/2013, 1 € = $ 1.3327
** Price for 1-year reserved instances/hour

Notes: a. No performance adjustment has been performed
b. Networking costs have been excluded in both cases
c. Storage costs have been excluded

PERFORMANCE DEGRADATION INCLUDED

e-FISCAL vs. Amazon EC2

e-FISCAL results compared with **EC2 reserved instances** (**all amounts in €**)

Costs refer to 2011 – Prices refer to 1/2013

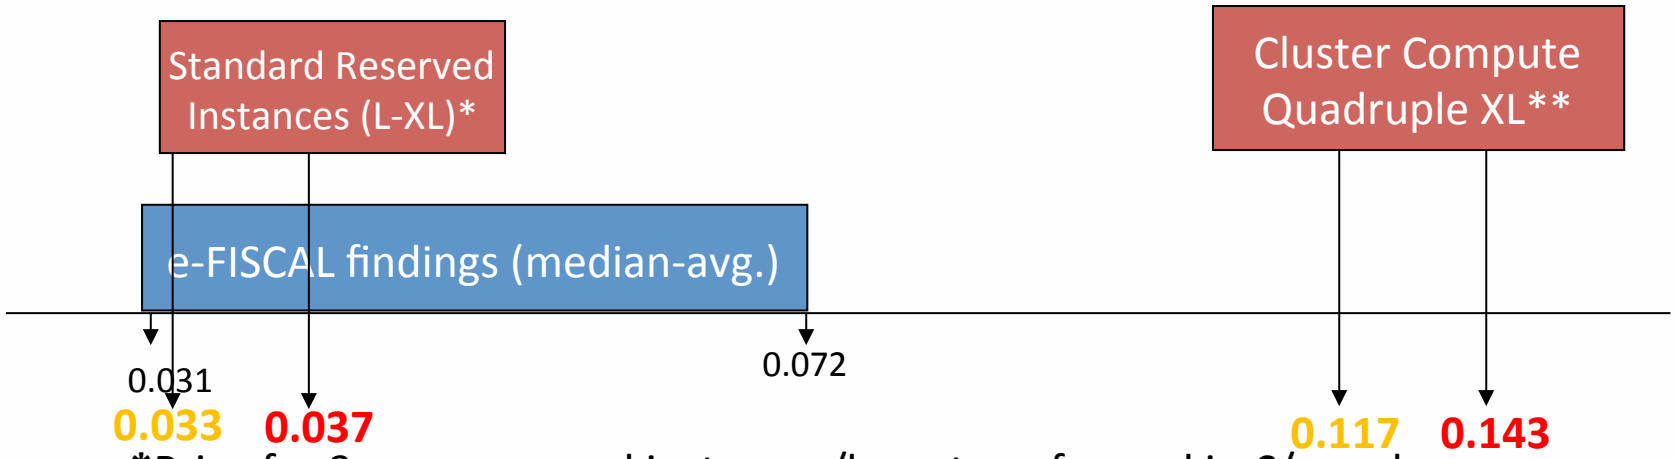

*Price for 3-year reserved instances/hour transformed in €/core hour equivalence based on instance characteristics)

Based on Linux/60% (red) -80% (yellow) usage of reserved instances.

Amazon site accessed on 15/1/2013, 1 € = $ 1.3327

** Price for 1-year reserved instances/hour

Notes: a. Performance adjustment has been performed
(Standard L-XL 27% / Cluster Compute Quadruple XL 43%)

b. Networking costs have been excluded in both cases

c. Storage costs have been excluded also

NO PERFORMANCE ADJUSTMENTS

e-FISCAL vs. Amazon EC2 (2)

e-FISCAL results compared with **EC2 on-demand instances** (**all amounts in €**)
Costs refer to 2011 – Prices refer to 1/2013

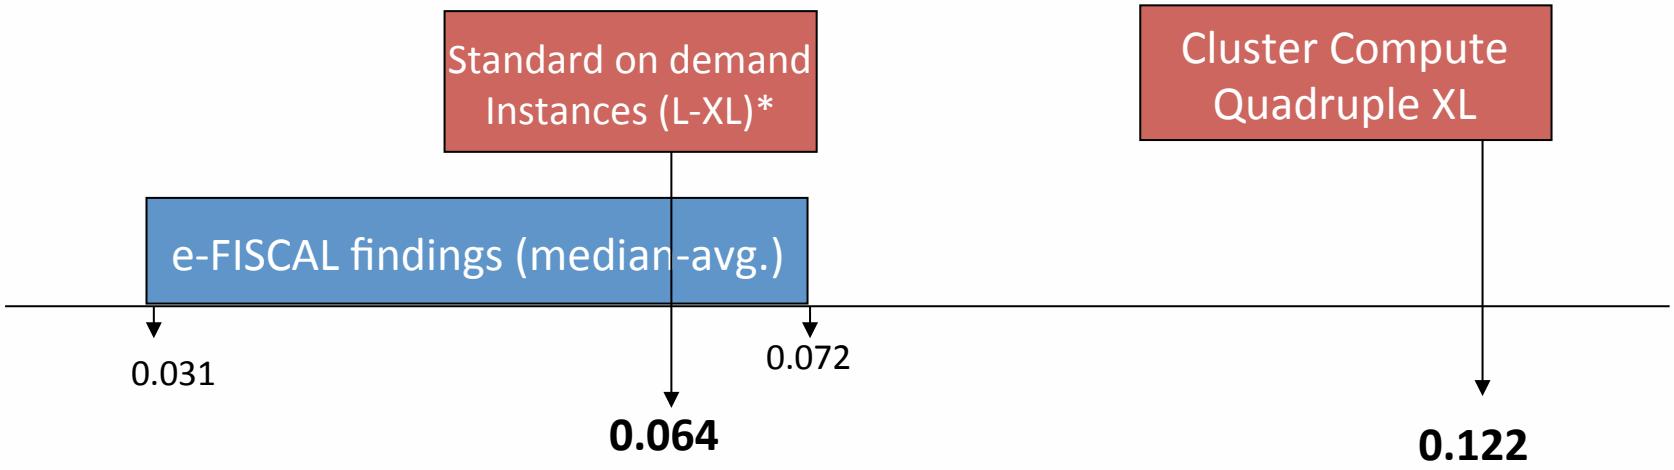

*Price for instances/hour transformed in €/core hour
(equivalence based on instance characteristics)
Based on Linux
Amazon site accessed on 15/1/2013, $ 1.3327

Notes: **a. No performance adjustment has been performed**
b. Networking costs have been excluded in both cases
c. Storage costs have been excluded also

PERFORMANCE DEGRADATION INCLUDED

e-FISCAL vs. Amazon EC2 (2)

e-FISCAL results compared with **EC2 on-demand instances** (**all amounts in €**)
Costs refer to 2011 – Prices refer to 1/2013

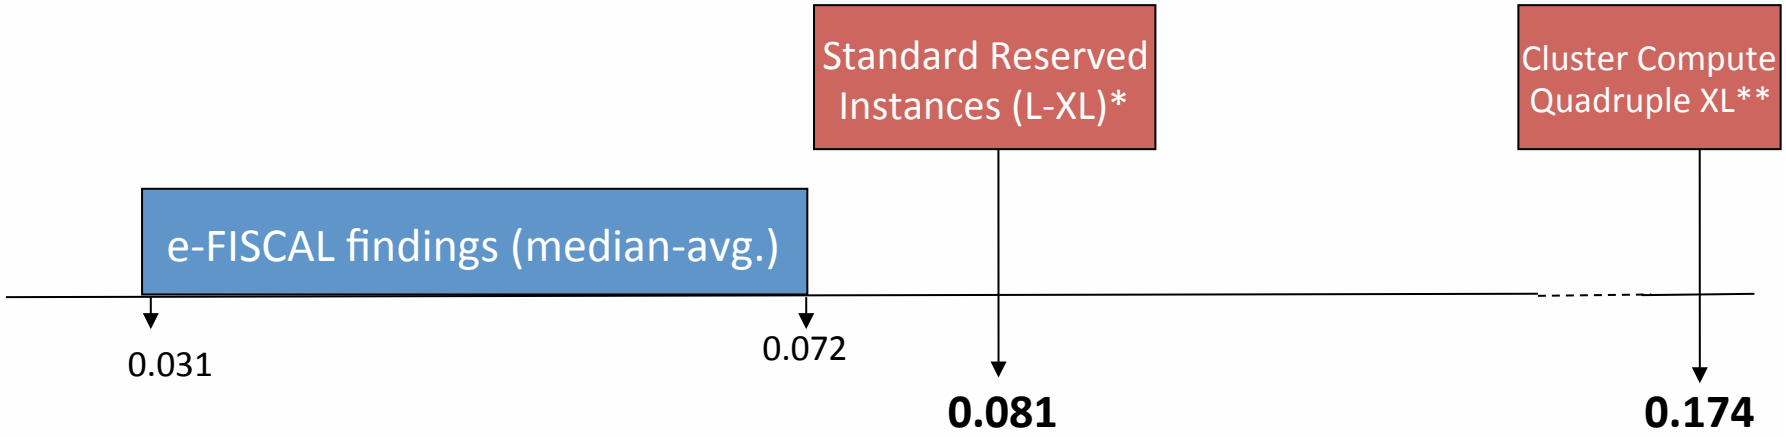

*Price for instances/hour transformed in €/core hour
(equivalence based on instance characteristics)
Based on Linux
Amazon site accessed on 15/1/2013, $ 1.3327

Notes: a. Performance adjustment has been performed
(Standard L-XL 27% / Cluster Compute Quadruple XL 43%)
b. Networking costs have been excluded in both cases
c. Storage costs have been excluded also

In-house utilization vs. Amazon (1)

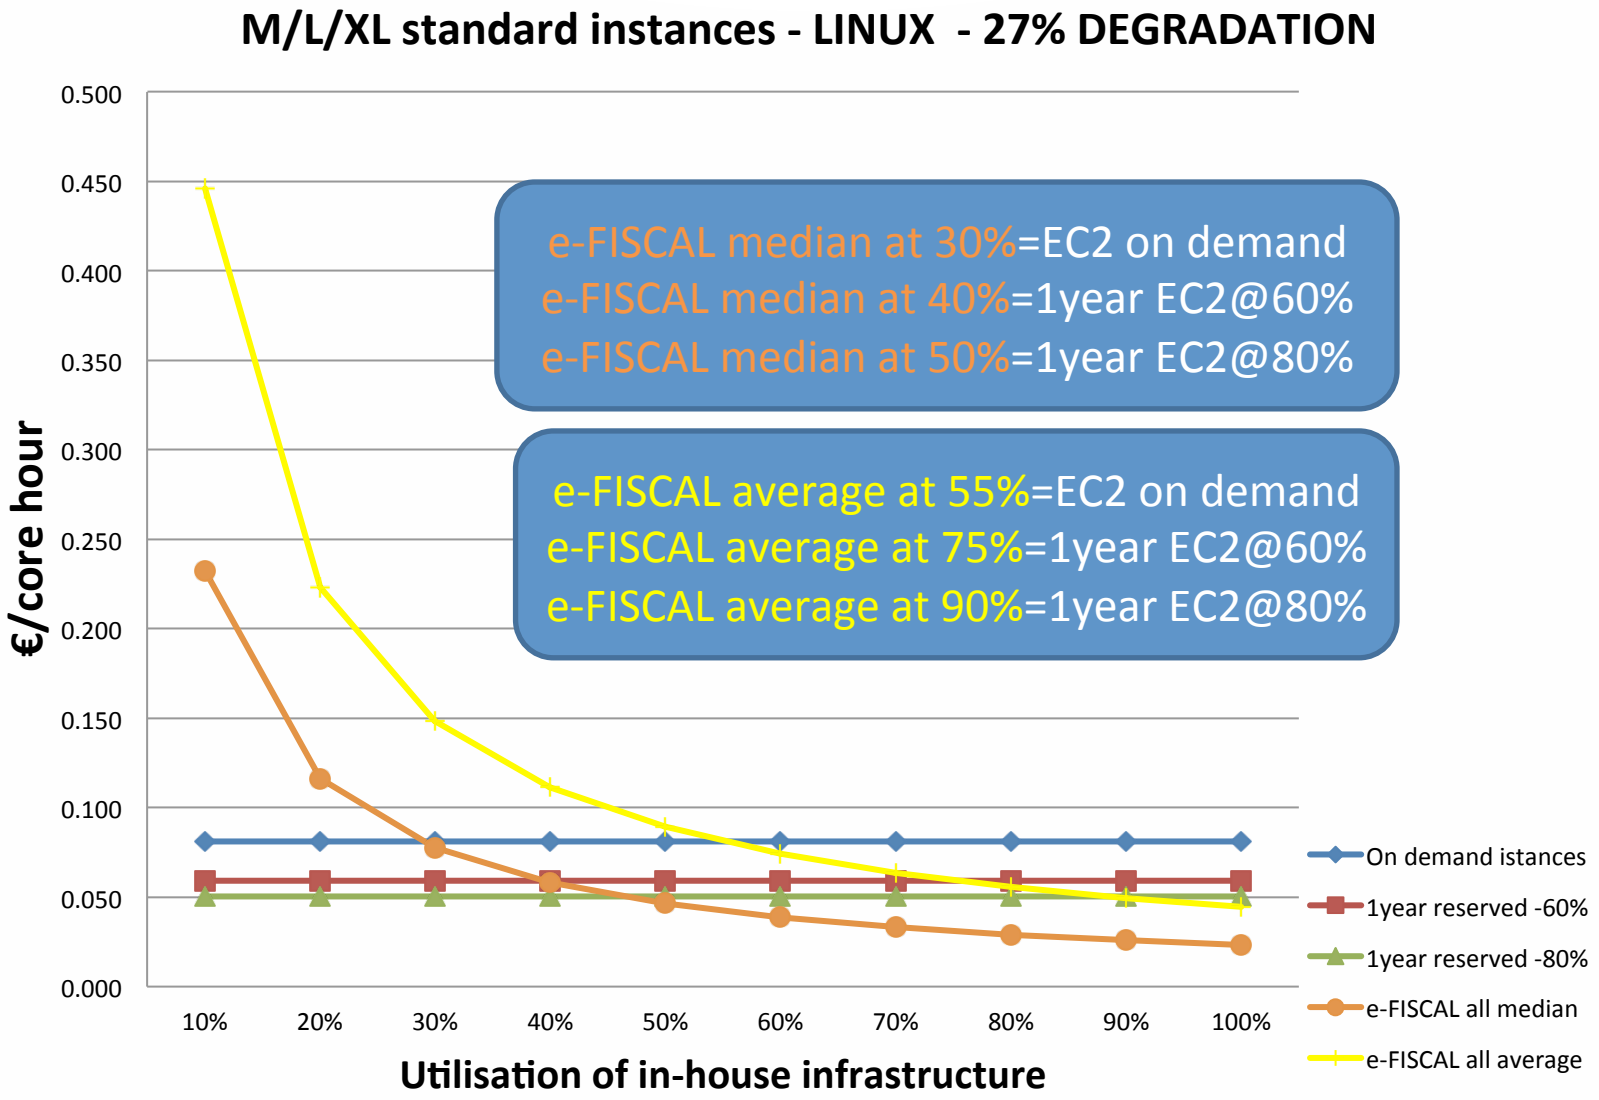

In-house utilization vs. Amazon (2)

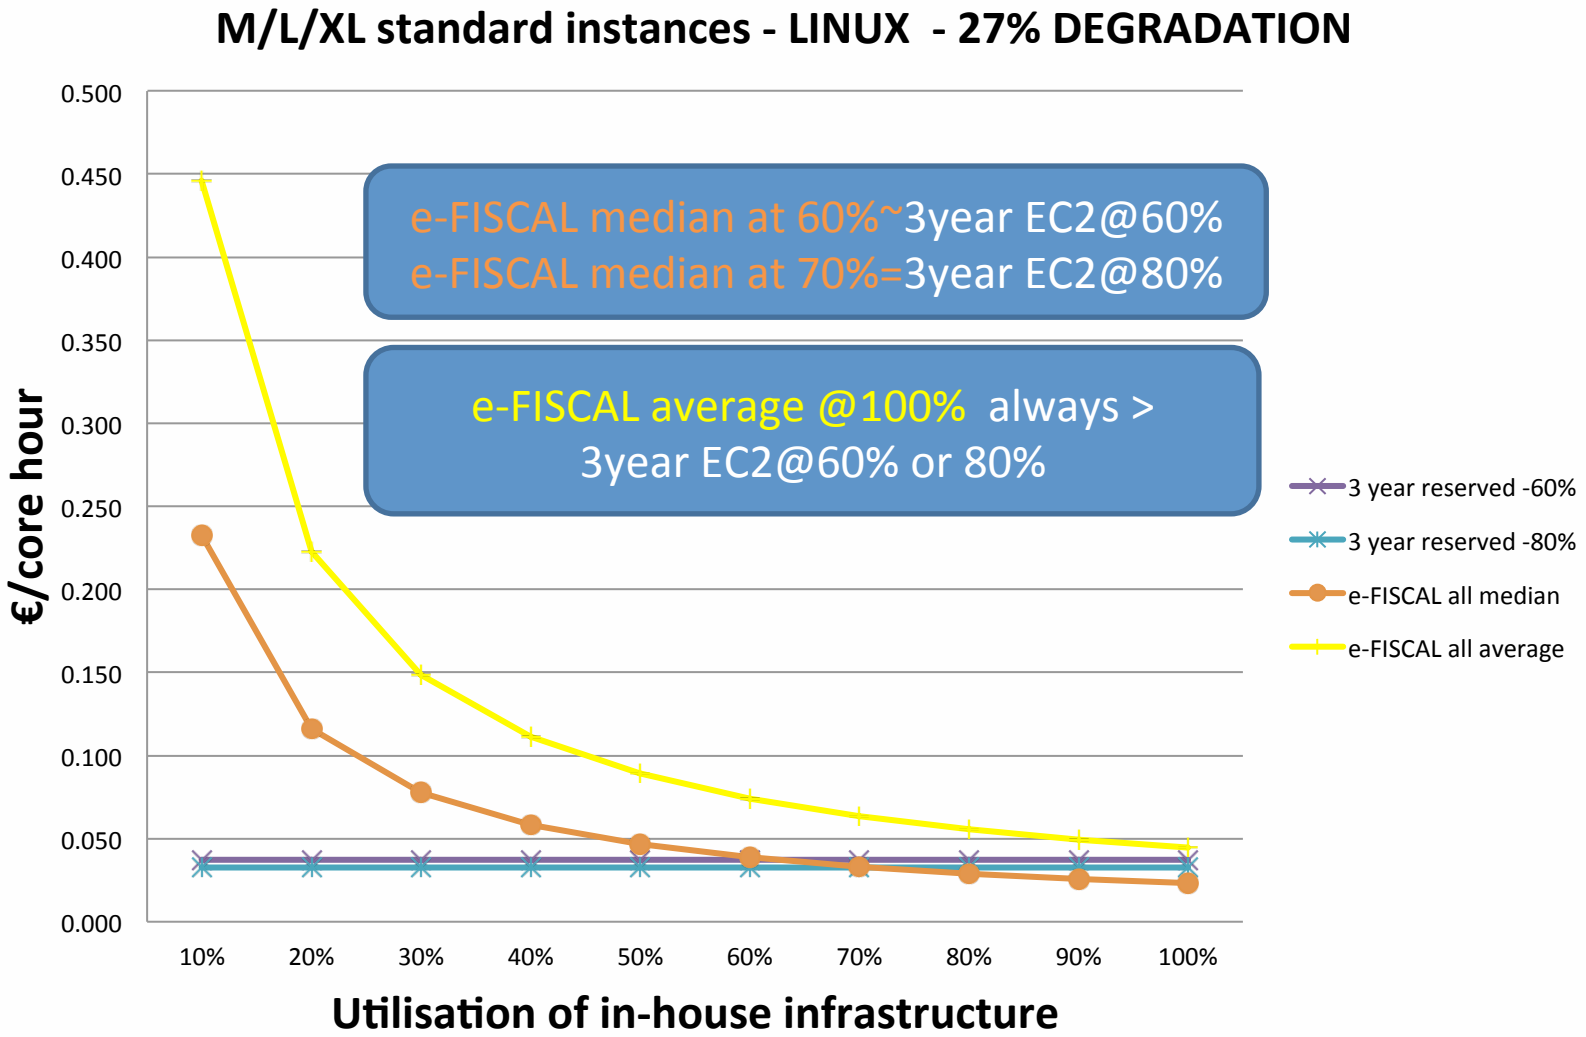

In-house utilization vs. Amazon (3)

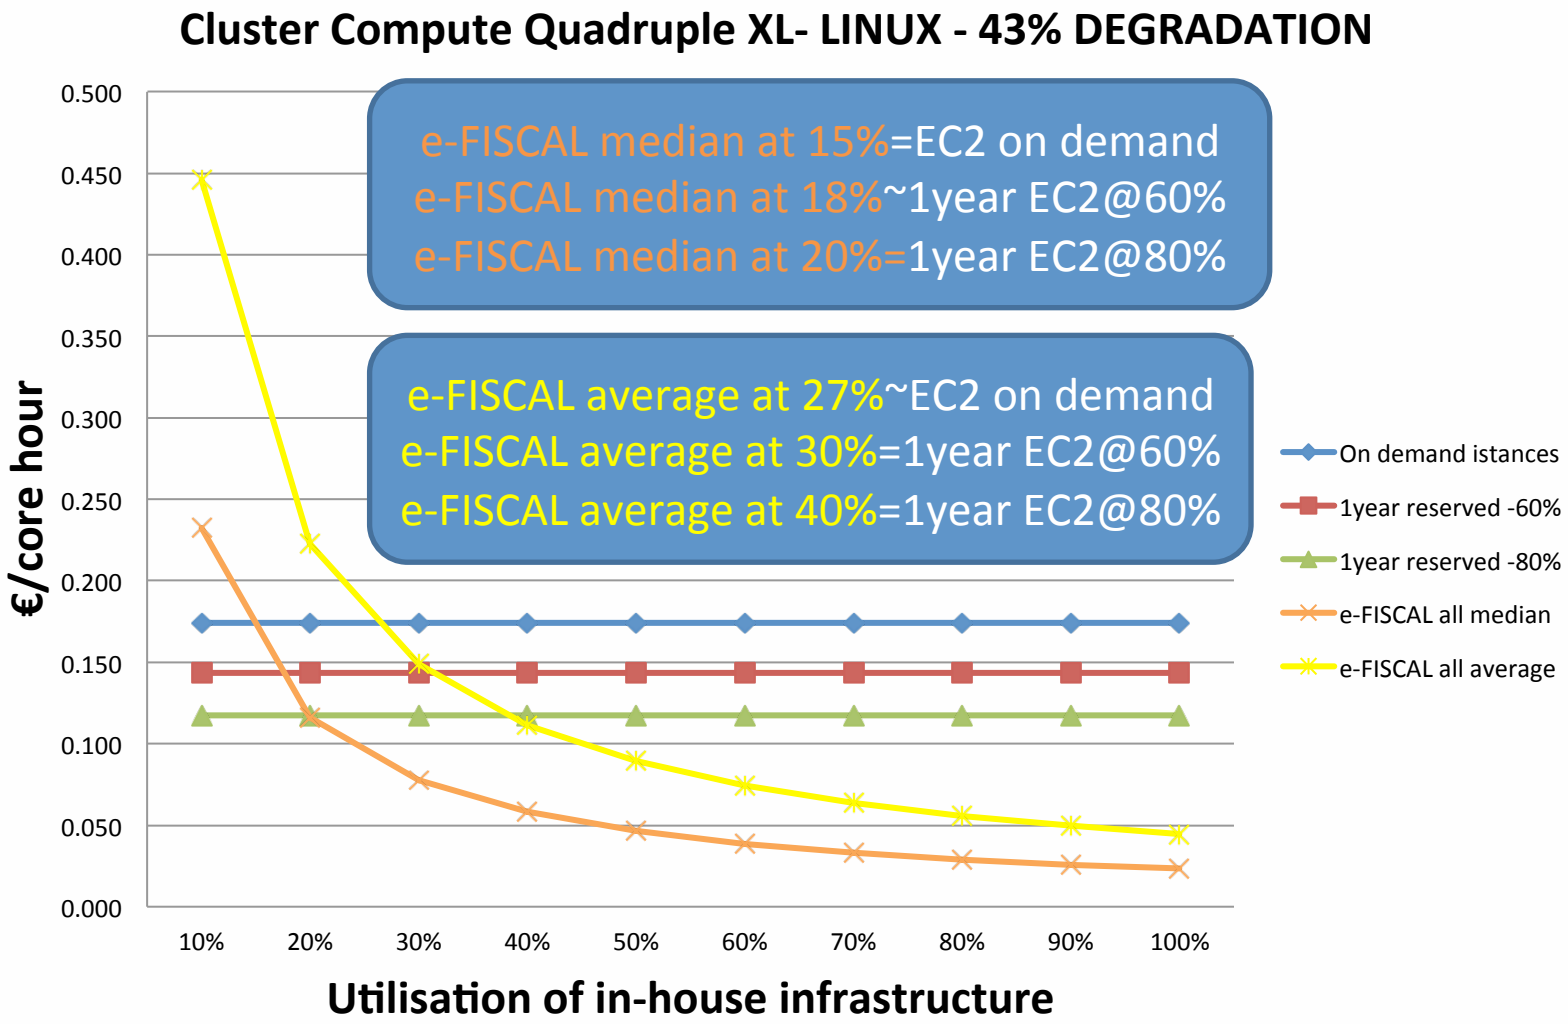

Conclusions

- **e-FISCAL pioneer in costing computing e-Infrastructures** :
  - Assessing costs in a highly distributed-heterogeneous environment!
- Our results are inline with literature
  - Cost per logical CPU/hour € 0.031 (median 2011 whole sample)
  - Costs show decreasing trends - not only for CAPEX but also for OPEX
- Nevertheless some interesting issues emerged:
  - Divergence in cost structures
  - High Useful lives
  - FTEs/core and personnel costs
  - Non- unanimous economies of scale existence
  - Moving to the cloud will not certainly reduce much on FTEs
  - Indifference points for in-house utilisation vary a lot:
    - 30-55% compared to EC2 standard on-demand instances, reserved ones competitive
    - 15-27% compared to cluster compute on demand, 18-40% for reserved ones

Thanks!

- All material to be available in www.efiscal.eu
  - e-mail us at info @ efiscal.eu to and keep up with the project (update list)

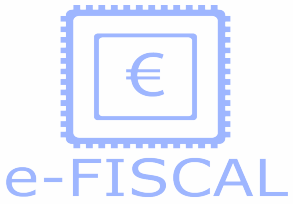

- **Project acronym:** e-FISCAL
- **Contract n°:** RI-283449
- **Project type:** CSA-SA
- **Start date:** 01/08/2011
- **Duration:** 18 months (end 31/1/2013)
- **Total budget:** 392.523 €
- **Funding from the EC:** 349 999 €
- **Total funded effort in PMs:** 33.75
- **Web site:** www.efiscal.eu

Transforming instances into number of cores

Type of Instances	Number of cores
M1 Medium	2
M1 Large	4
M1 Extra Large	8
M3 Extra Large	13
M3 Double X Large	26
Cluster Compute	
Quadruple Extra Large	8
Eight X-Large	16
High-CPU Instances	
Extra Large	20

Sources: Berriman, B. and Deelman, E. "How To Use Cloud Computing To Do Astronomy", IPAC, May 9, 2012, p. 8; plus e-FISCAL estimations

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. Note: P

Answer Options	Min	Max	Average	Median	Answered questions
Cost per logical CPU in € in 2010	100	800	299	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	800	277	210	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

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

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%	7%	5%
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

Personnel costs - 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	50.58	44.55
Average yearly salary cost per FTE (gross salary plus employee benefits and bonuses) **in '000 €** in 2011	15	108	51.41	46.30

The salary range is very wide

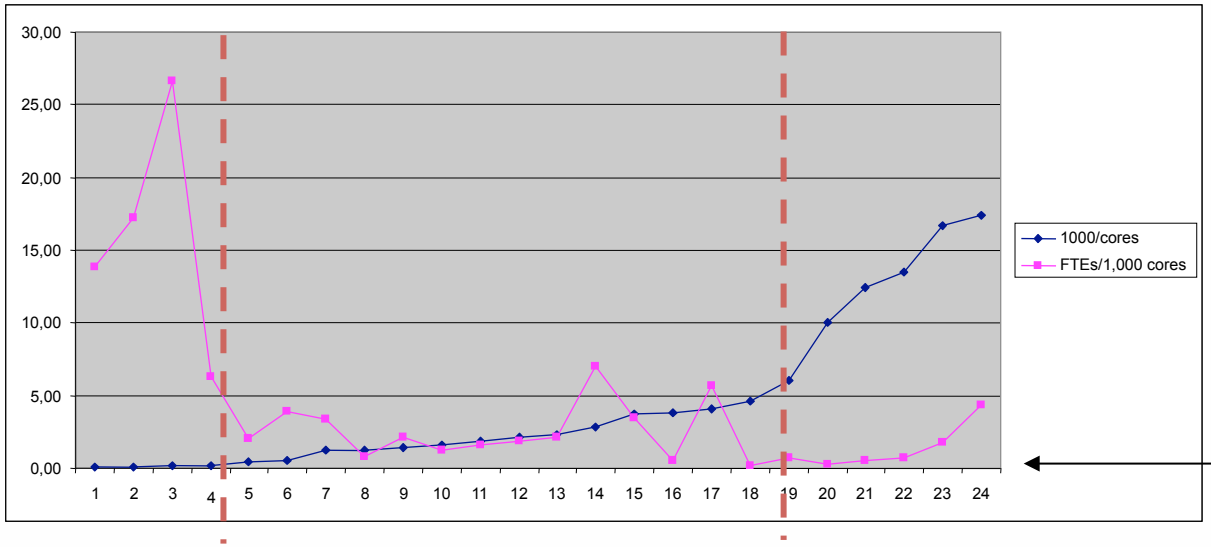

Plotting 1,000 Logical CPUs and number of FTEs per 1,000 Logical CPUs
Generally, no of FTEs/1,000 cores decreases as site size increases

Questionnaires 1 to 24

Power Usage Effectiveness

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	1.25	2.2	1.58	1.50
Power Usage Effectiveness in 2011	1.25	2.24	1.55	1.49

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.
- Appling 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