e-FISCAL

e-infrastructure CAPACITIES

e-FISCAL Workshop @ EGI TF 12 -
Prague, 21 September 2012

Introduction and key findings

Fotis Karagiannis, Sandra Cohen,

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

It's all about knowing the costs..

…their composition..

…and putting them in context!

e-FISCAL

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*

e-FISCAL

Background

- First in-depth study at European scale
  - Significant sample of participants, HTC/HPC, comparisons with Clouds
- Builds on previous financial exercise
  - e-IRGSP2 project
  - Dealt with HTC (Grids) only, small number of NGIs involved -> initial charting of the area
  - Findings available at http://www.e-irg.eu/images/stories/e-irgsp2_d4_3_approved_by_the_consortium.pdf (look at deliverable second part)

Basis of costing exercise

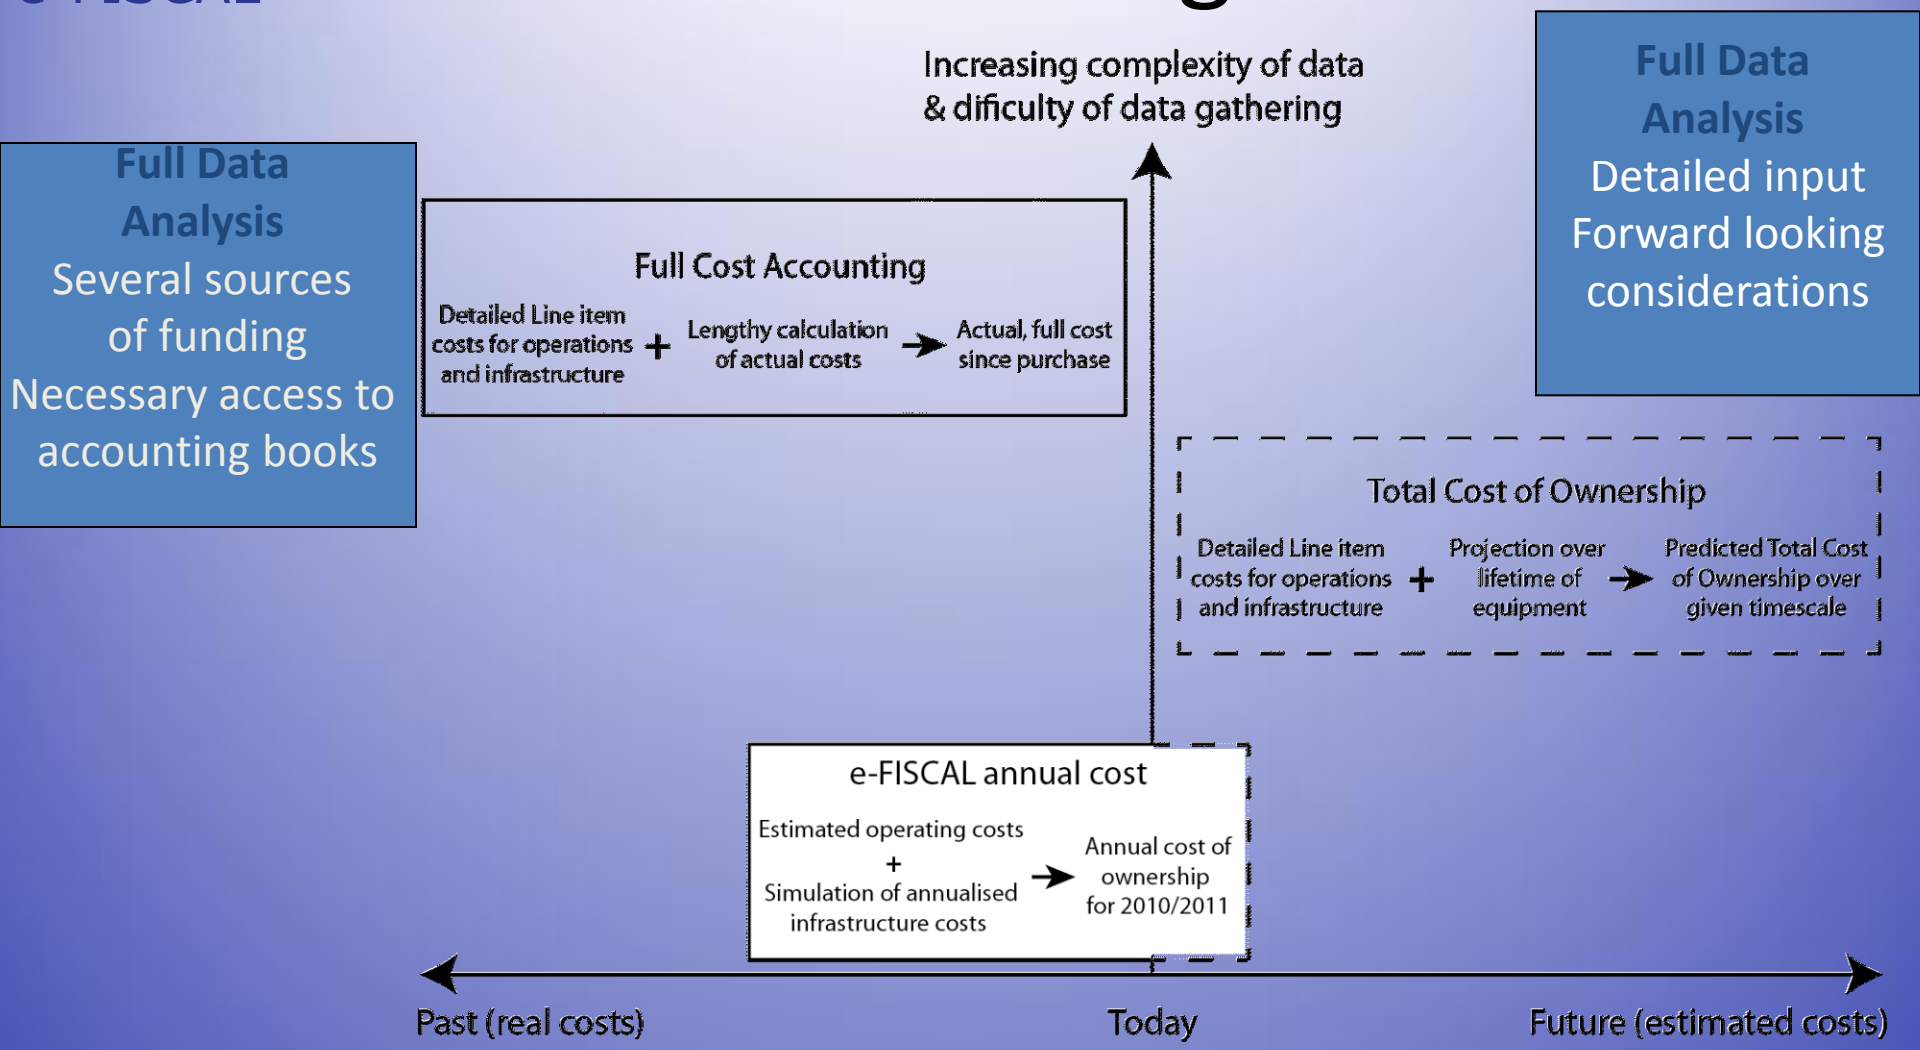

Methodology overview

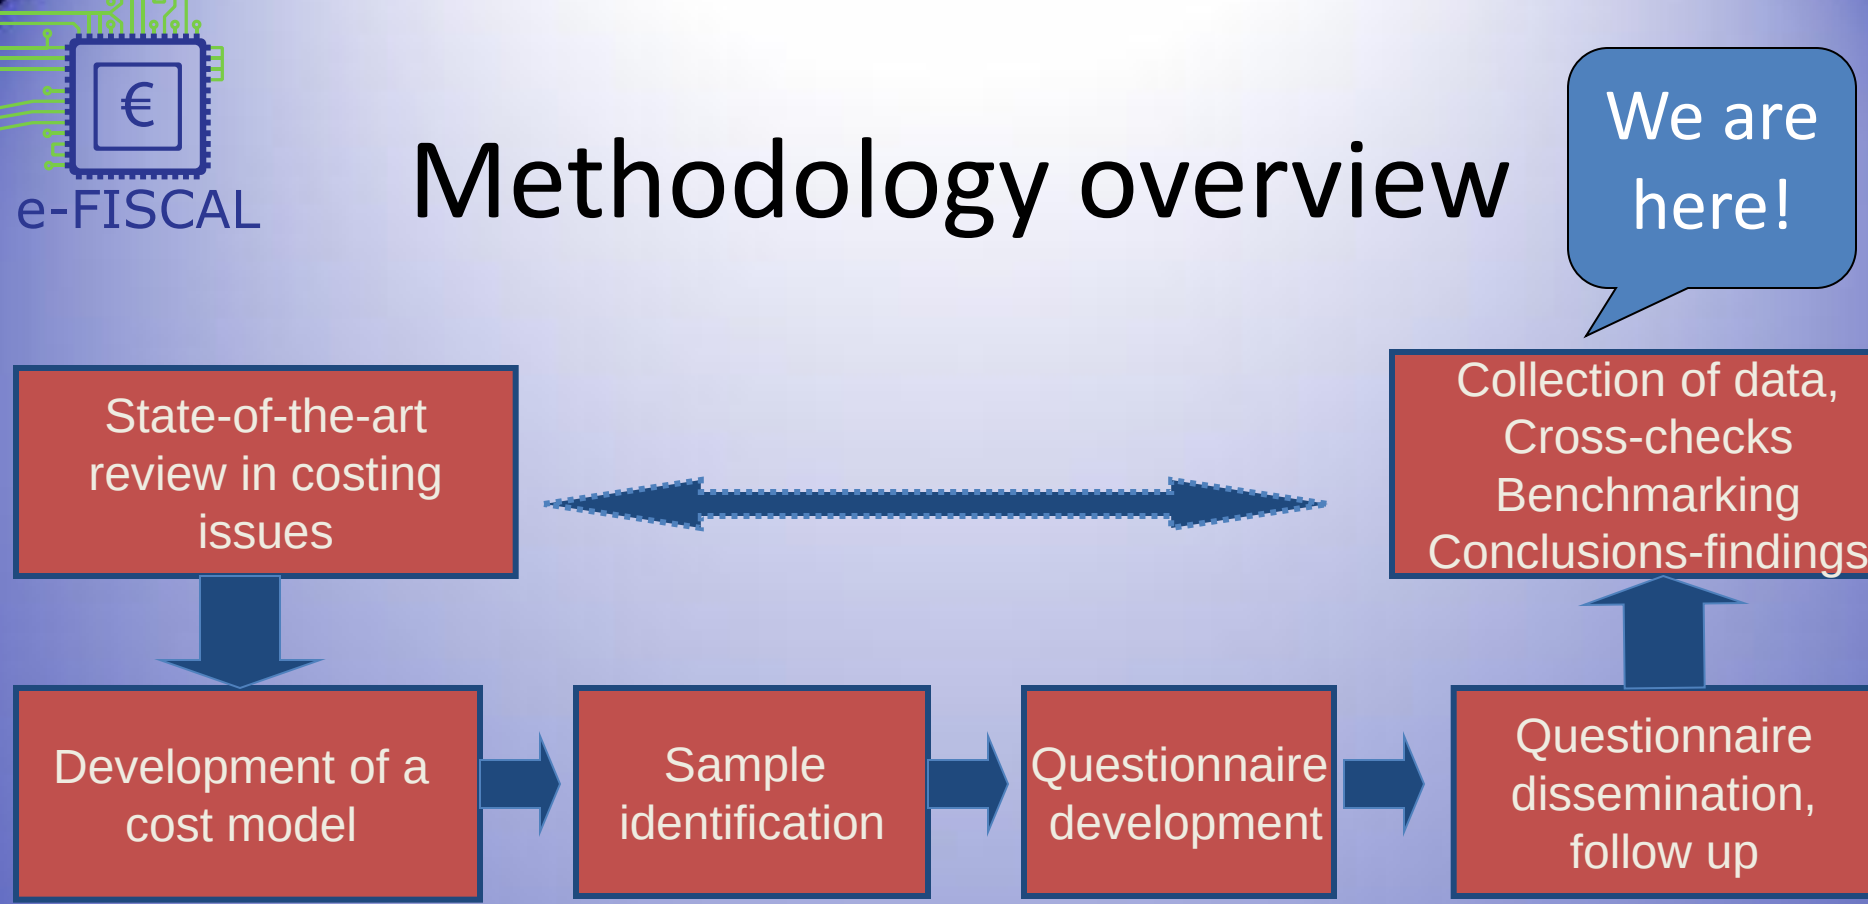

We have gone through the first full cycle of the methodology and we are about to start again by capitalizing on the feedback and experience gained

Contributions/disclaimers

- Disclaimers:
  - Careful in comparing e-Infrastructure costs with Cloud prices!
    - benchmarking,
    - profit margin possible
    - however a user cares about the actual cost
  - Confidentiality/Anonymity of data!
    - Cross-checks/validation with market or other prices
    - **No identifiable data related to an individual site or national HPC/HTC entity are presented**
- Cost is different from value!

Countries contributing

The study is on-going

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

Sample/Respondents so far…

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

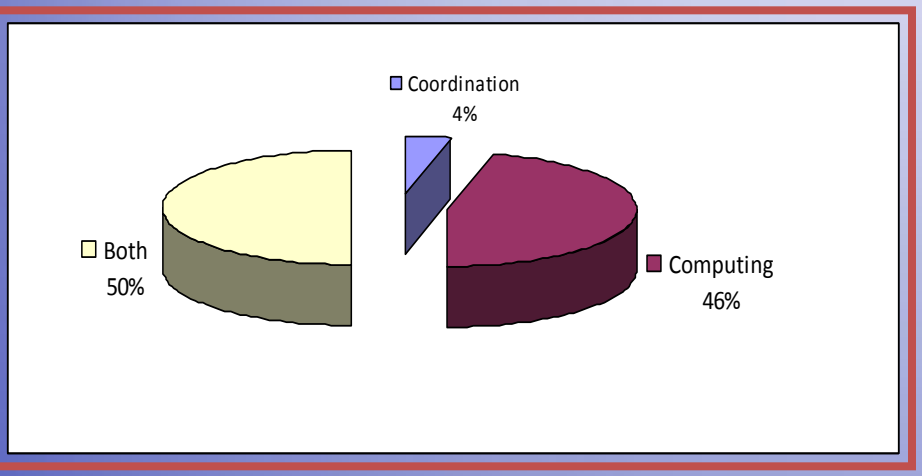

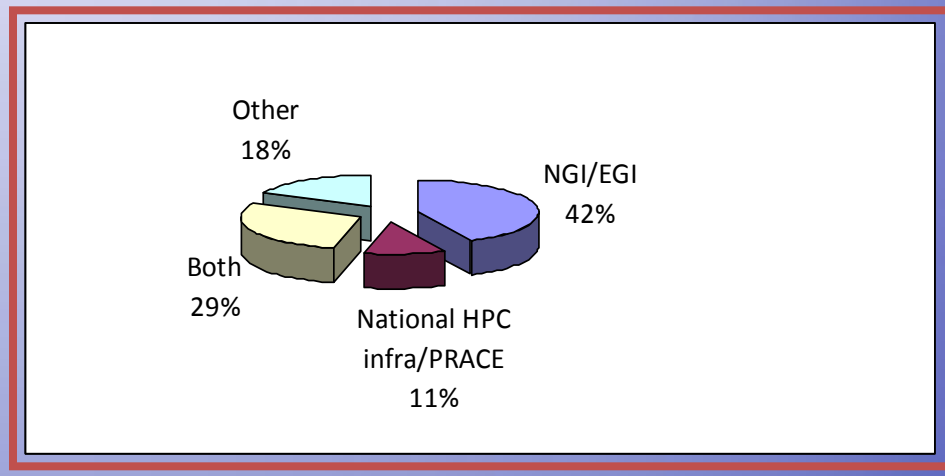

  - The vast majority of respondents provide both computing and coordination
  - Most of the data from HTC or mixed HTC/HPC centres

10

Review the state-of-th

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)

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

e-FISCAL: first conclusions

- e-FISCAL results in-line with the literature
- In-house HPC/HTC e-Infrastructures are cost-effective (w. high utilisation rates & depreciation rates)
  - however use case-based analysis important!
- Personnel ~50% of total costs; CAPEX/OPEX=30/70%
- Larger sites have in general less FTEs/core and lower cost per core hour
- Initial (small-scale) benchmarking efforts between in-house HPC and Amazon Compute Cluster instance:
  - A ~40% performance degradation of the latter for HPC, similar for HTC
- 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:	0,073	-	**0,031**

Median for minimum utilisation rate: 74%

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:	277		**210**
• Average 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:	51k		**46k**
• Power Usage Effectiveness:	1,55		**1,49**

Costs breakdown (2011-median)

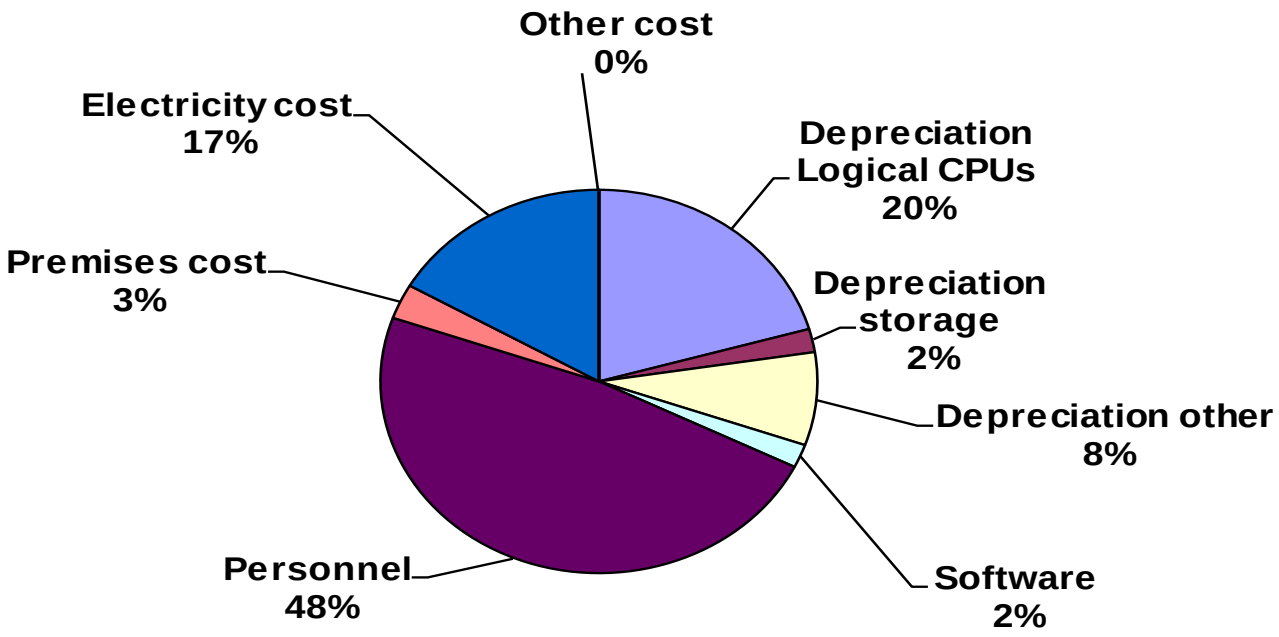

Cost per core hour in € / no of cores*

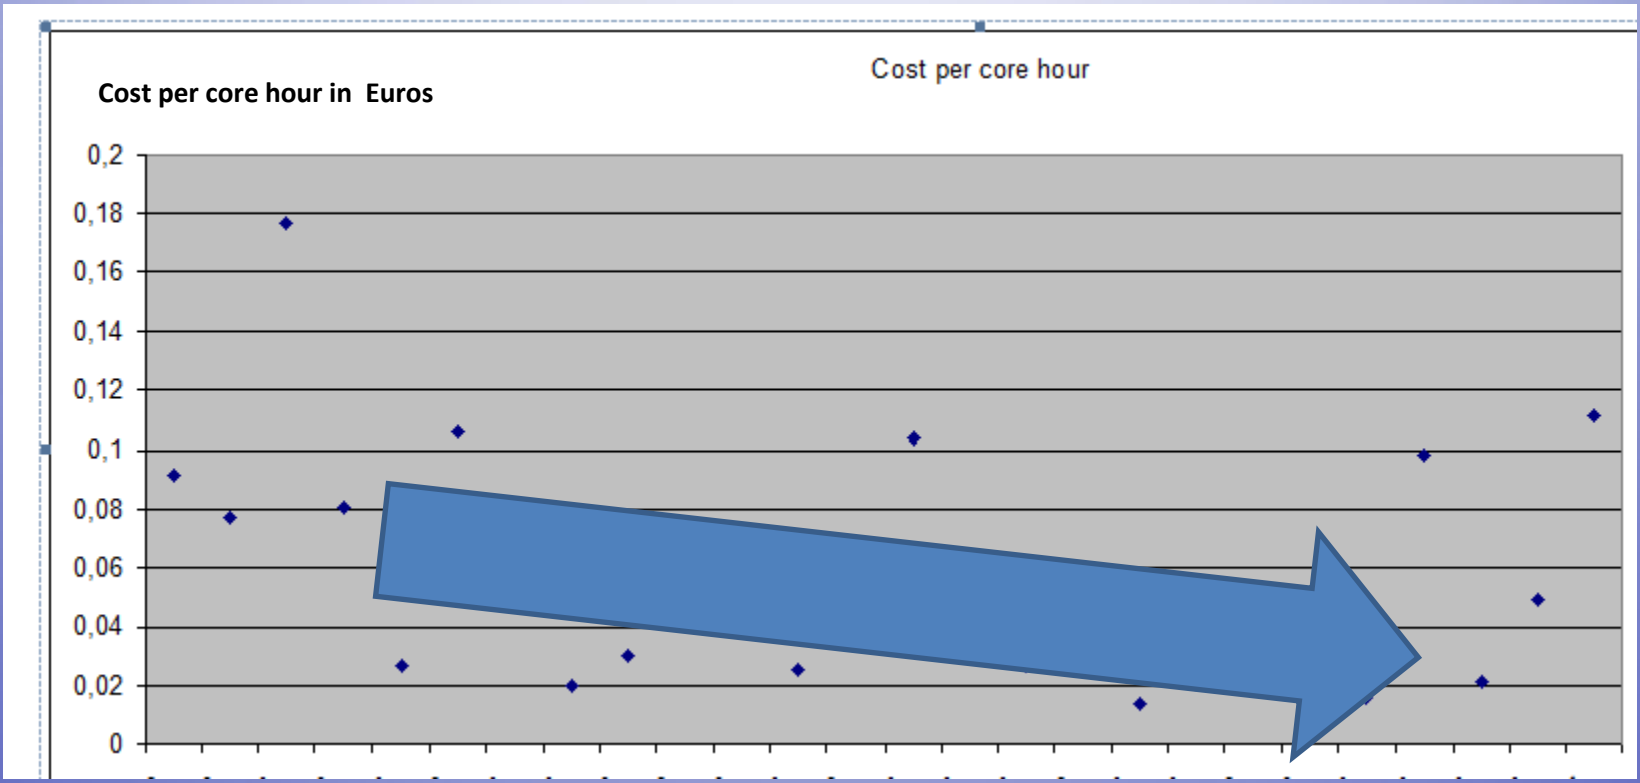

* Dots are sites!

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

e-FISCAL vs. Amazon EC2

e-FISCAL results compared with **EC2 reserved instances** as (**all amounts in €**)
Costs refer to 2011 – Prices refer to 9/2012

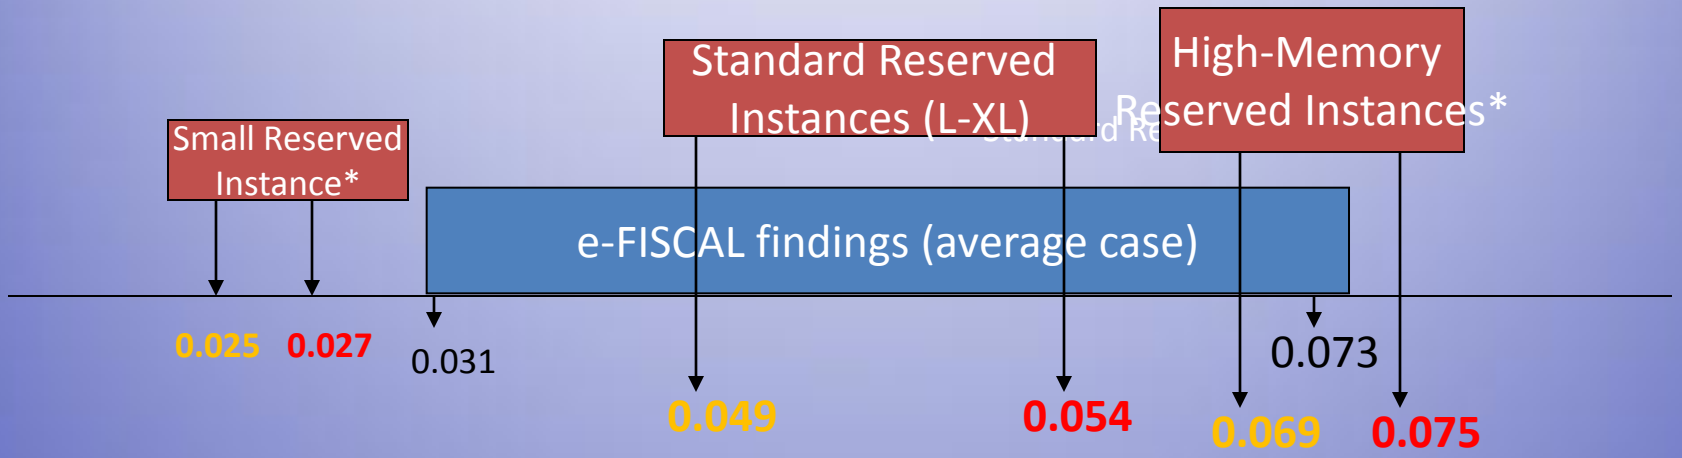

*Cost for 3-year reserved instances/hour
transformed in €/logical CPU hour (equivalence based on instance characteristics)
Based on windows/EU-Ireland/80% (red) -100% (yellow) usage of reserved instances.
Amazon site accessed on 12/9/2012, 1 € = $ 1,2878

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

e-FISCAL vs. Amazon EC2 (2)

e-FISCAL results compared with **EC2 on-demand instances** as (**all amounts in €**)
Costs refer to 2011 – Prices refer to 9/2012

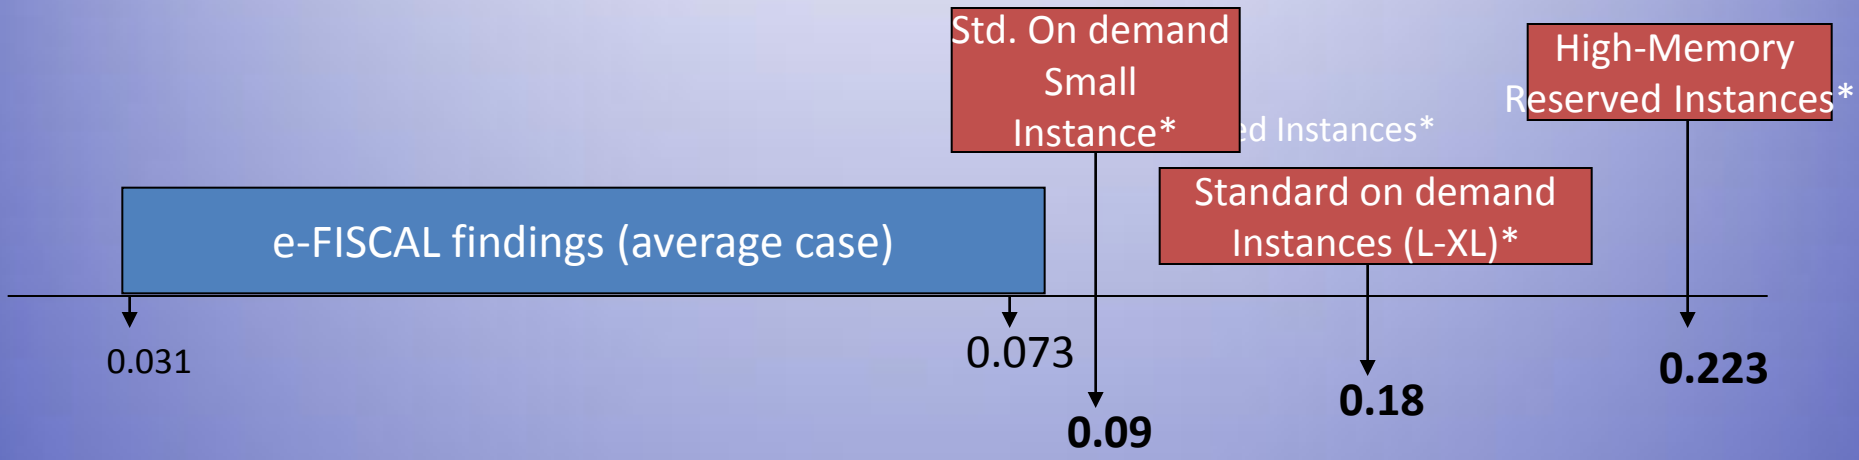

**Cost for instances/hour
transformed in €/logical CPU hour (equivalence based on instance characteristics)
Based on windows/EU-Ireland
Amazon site accessed on 12/9/2012, 1 € = $ 1,2878

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

Transforming instances into number of cores

	Number of cores
Standard Instances	
Small (Default)	**1**
Large	**2**
Extra Large	**4**
High-Memory Instances	
Extra Large	**2**
Double Extra Large	**4**
Quadruple Extra Large	**8**

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

Conclusions

- **e-FISCAL novelty**: Assessing and comparing 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
      - Evidence of existence of economies of scale
- Nevertheless some interesting issues emerged:
  - Divergence in cost structures
  - High Useful lives
  - FTEs/core 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!

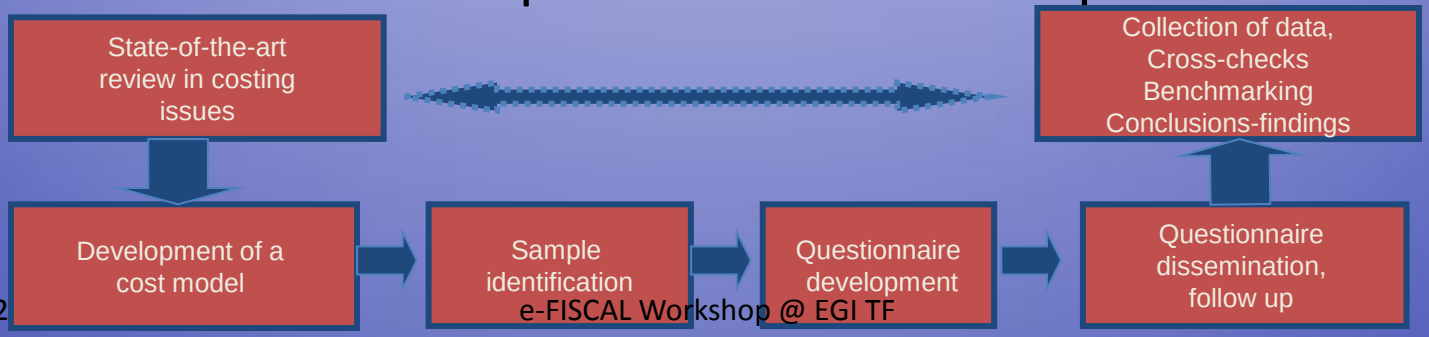

Thanks!

- 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

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

23

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

24

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

25

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

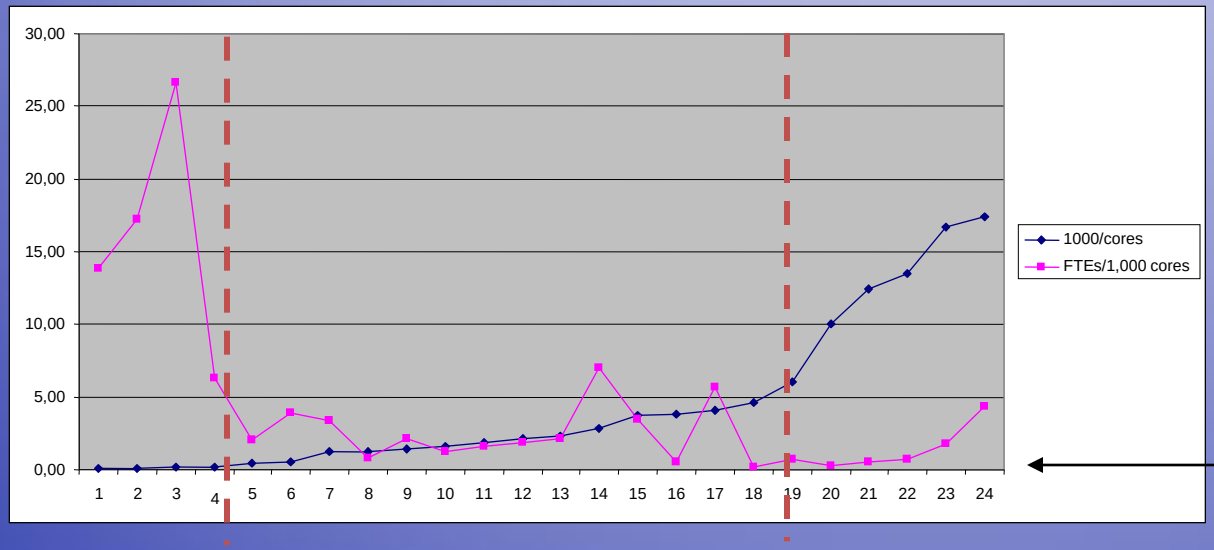

Plotting 1,000 Logical CPUs and number of FTEs per 1,000 Logical CPUs
Generally, no of FTEs/1,000 cores decreases as site size increases

Questionnaires 1 to 24

26

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

27