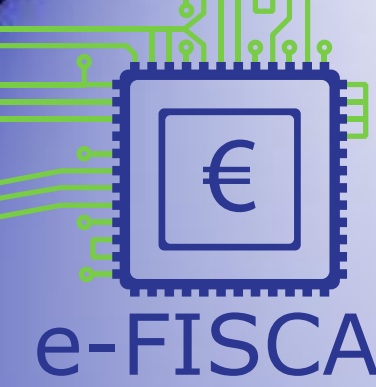

HelixNebula Workshop @ EGI TF 12 -
Business Models and Legal Aspects Session
Prague, 19 September 2012

The Cost of Computing e-Infrastructures

Fotis Karagiannis,

Independent/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

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

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

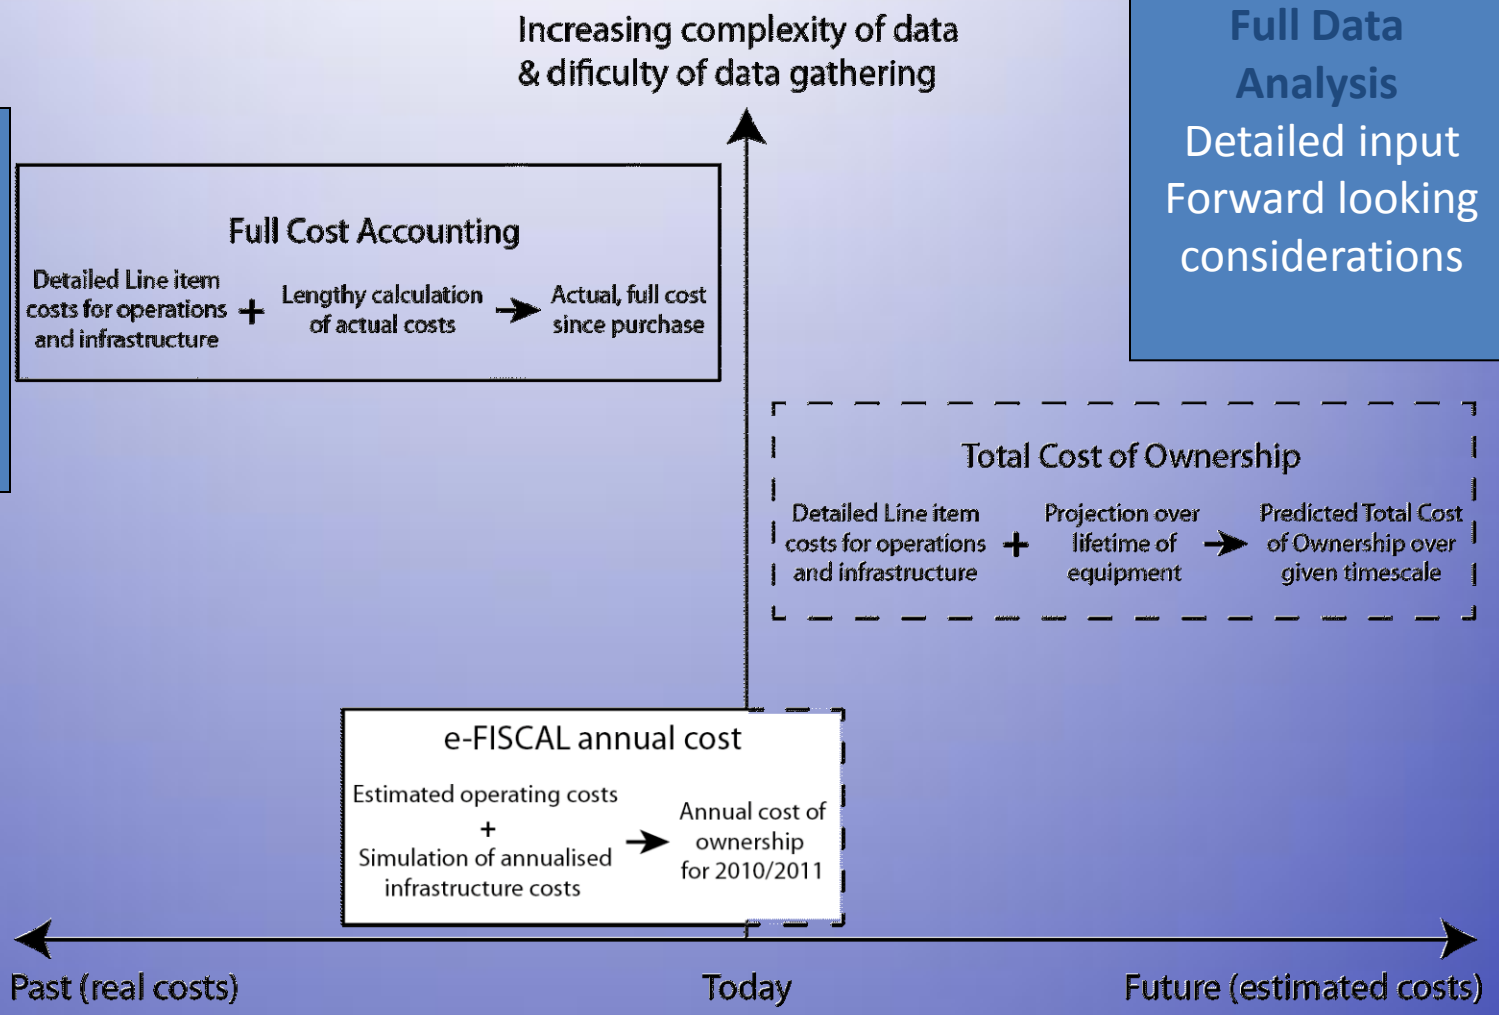

Full Data Analysis
Detailed input
Forward looking considerations

Methodology overview

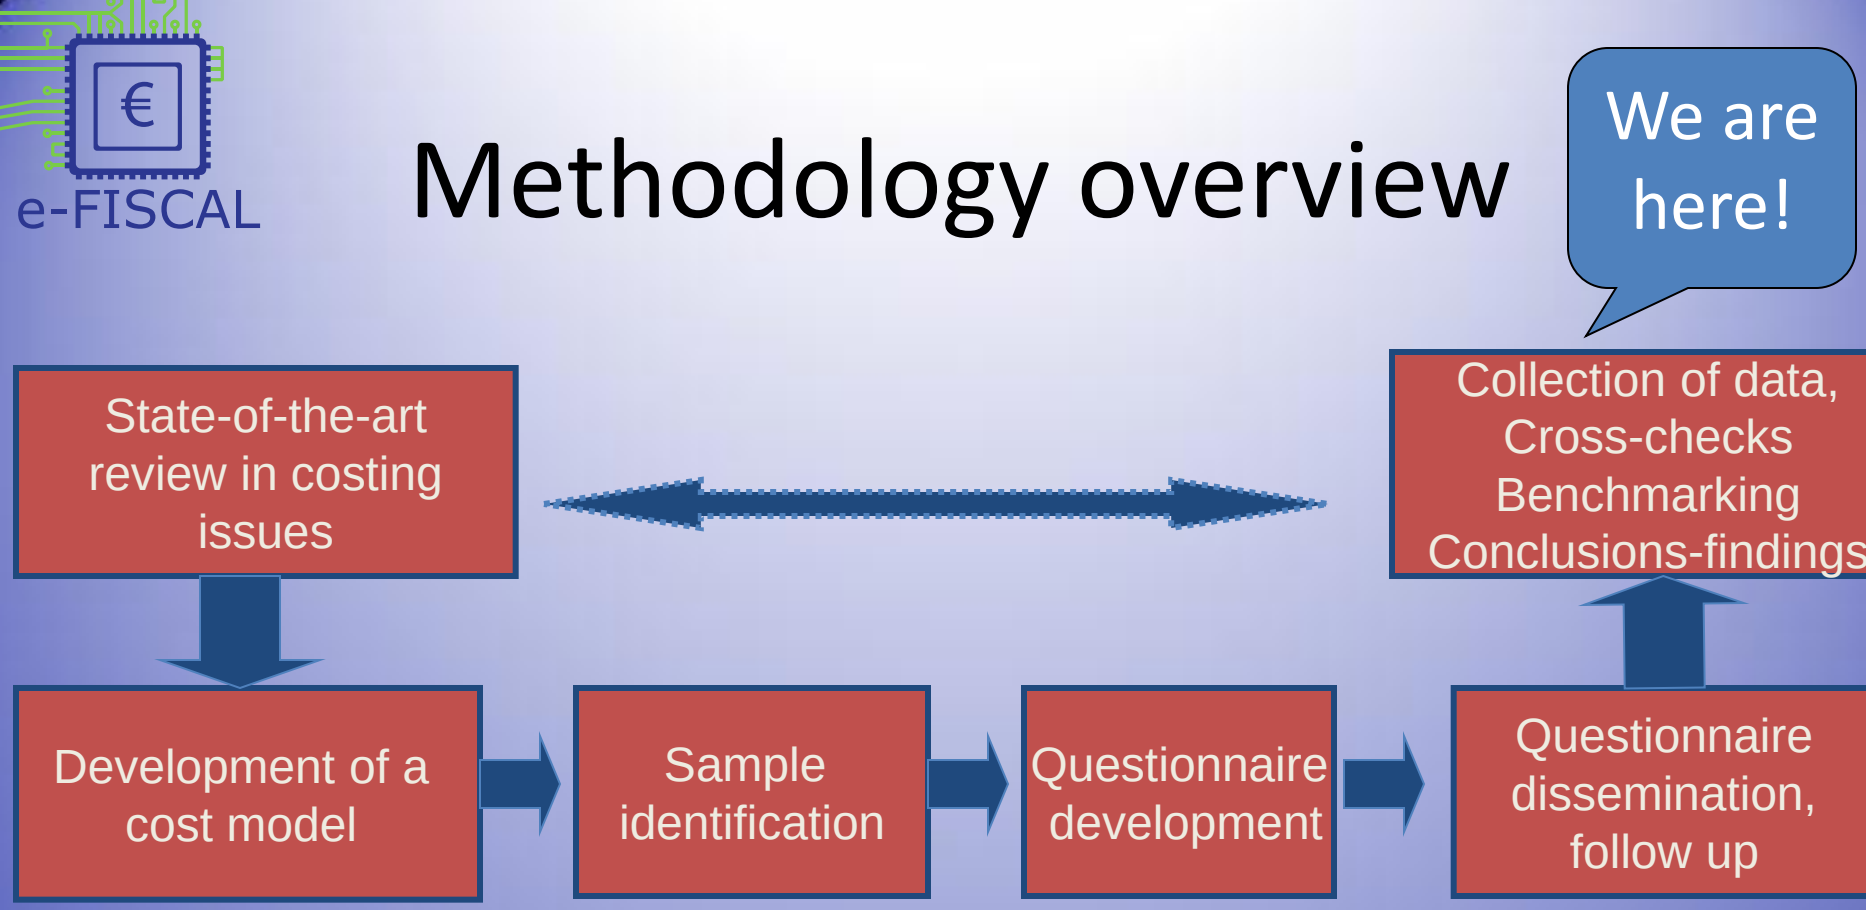

We have gone through the first full cycle of the methodology and we are about to start again by capitalizing on the feedback and experience gained

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

Review the state-of-the-art

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**

Cost per core hour in € / no of cores*

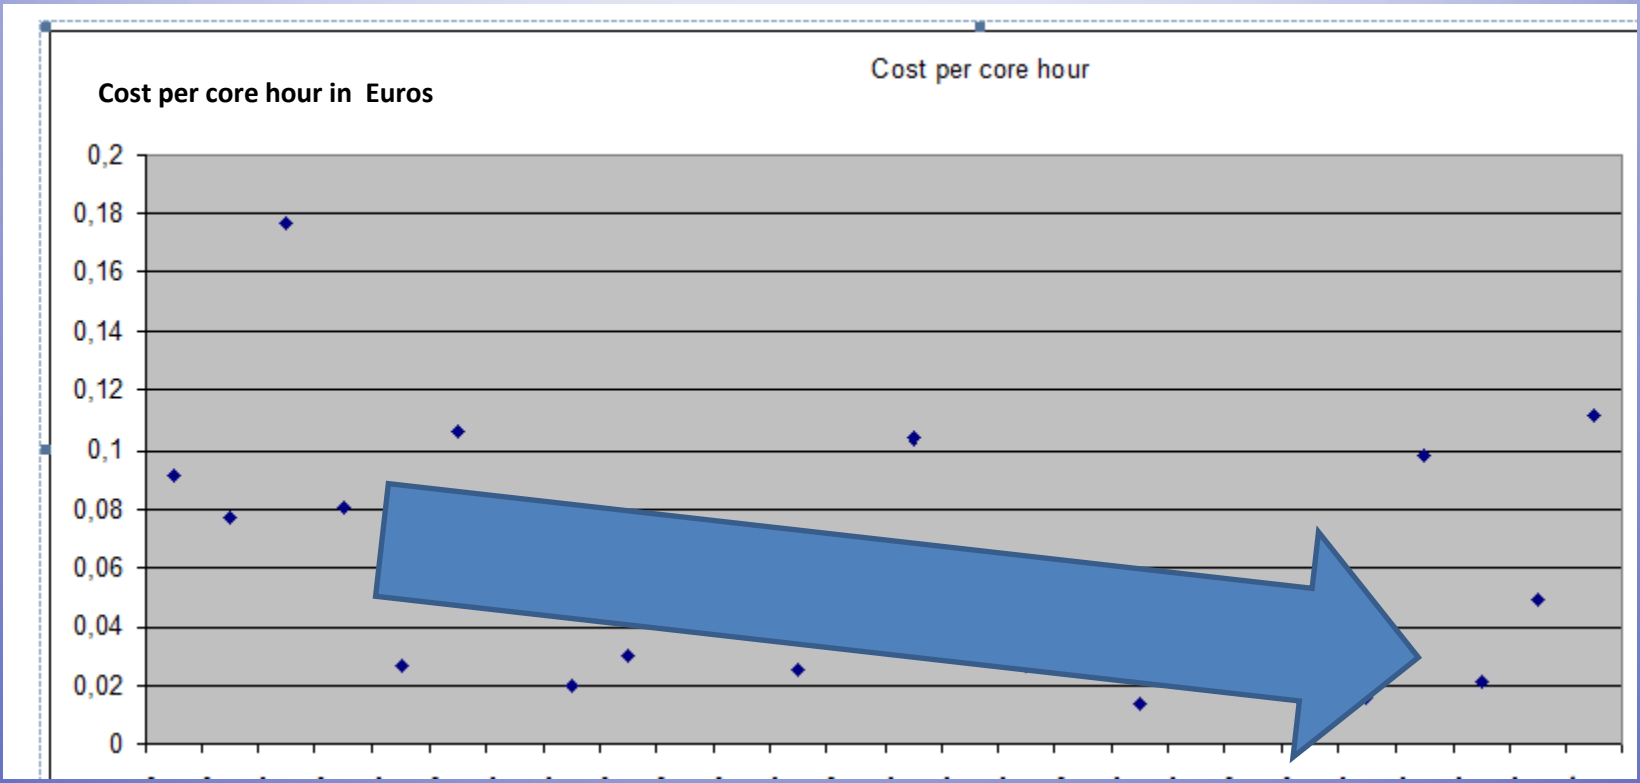

* Dots are sites!

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

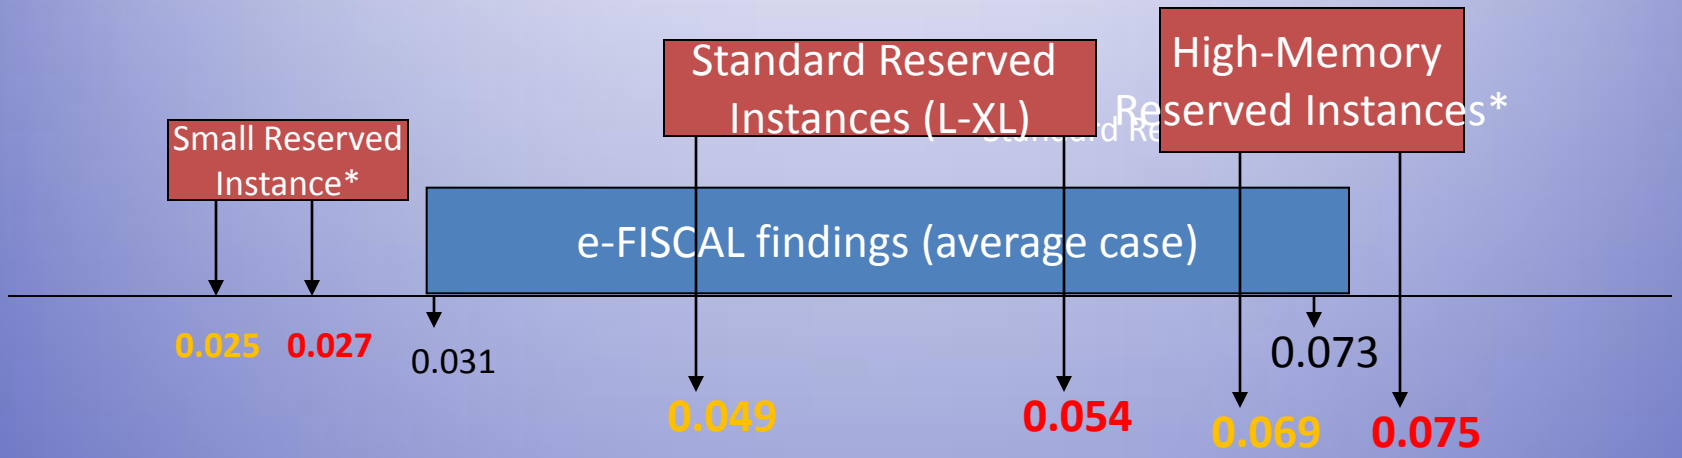

*Cost for 3-year reserved instances/hour
transformed in €/logical CPU hour (equivalence based on instance characteristics)
Based on windows/EU-Ireland/80% (red) -100% (yellow) usage of reserved instances.
Amazon site accessed on 12/9/2012, 1 € = $ 1,2878

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

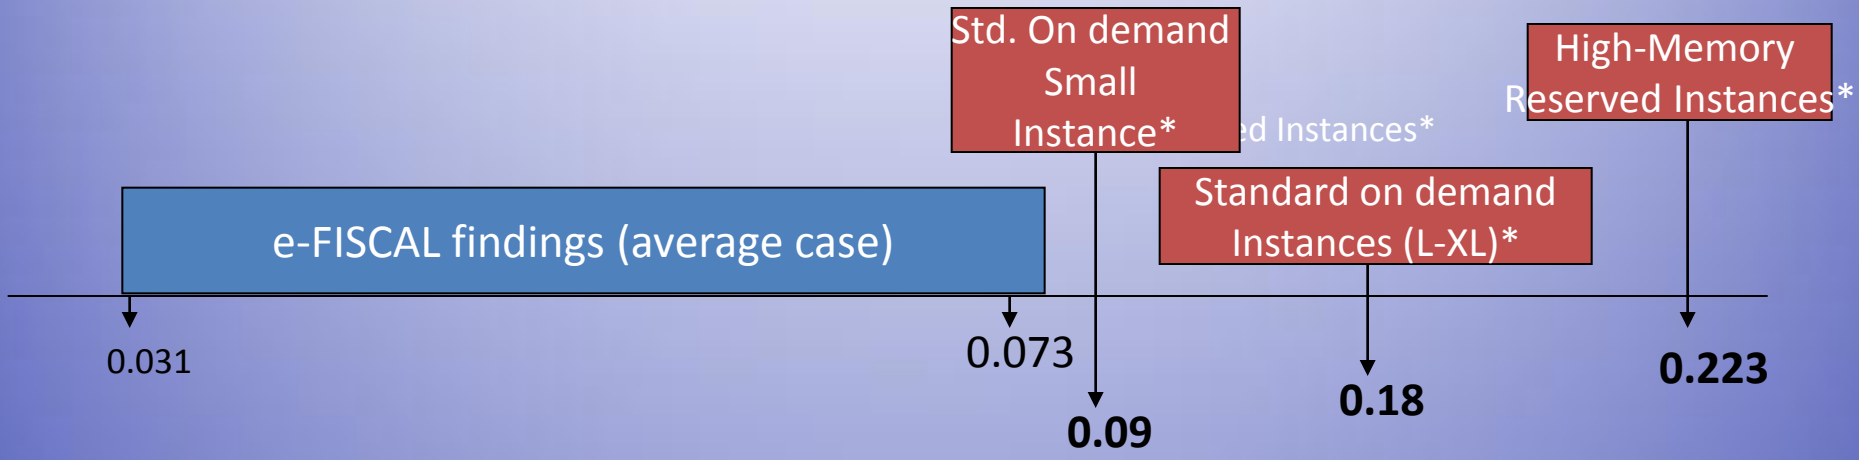

**Cost for instances/hour
transformed in €/logical CPU hour (equivalence based on instance characteristics)
Based on windows/EU-Ireland
Amazon site accessed on 12/9/2012, 1 € = $ 1,2878

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!

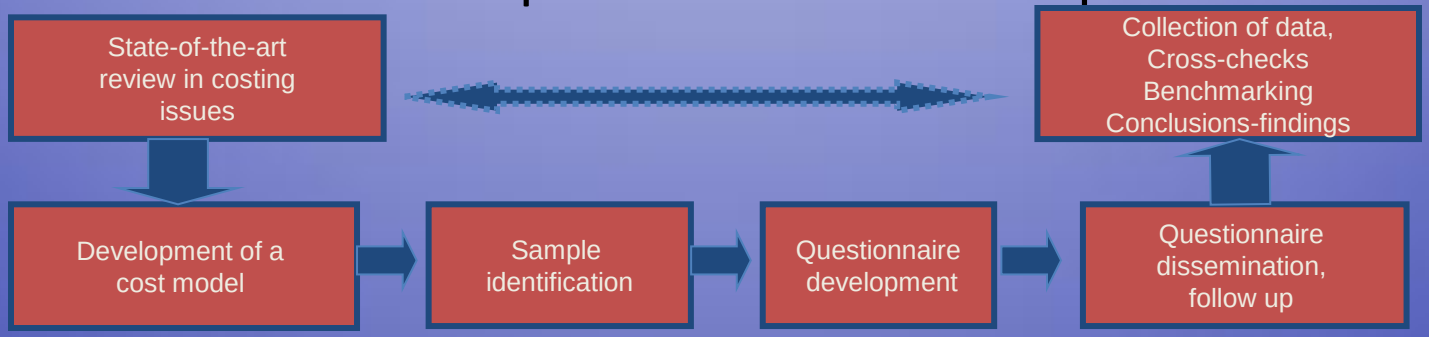

Thanks!

- 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)

e-FISCAL

e-FISCAL

Project in a nutshell

- **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

Business models for Research (1)

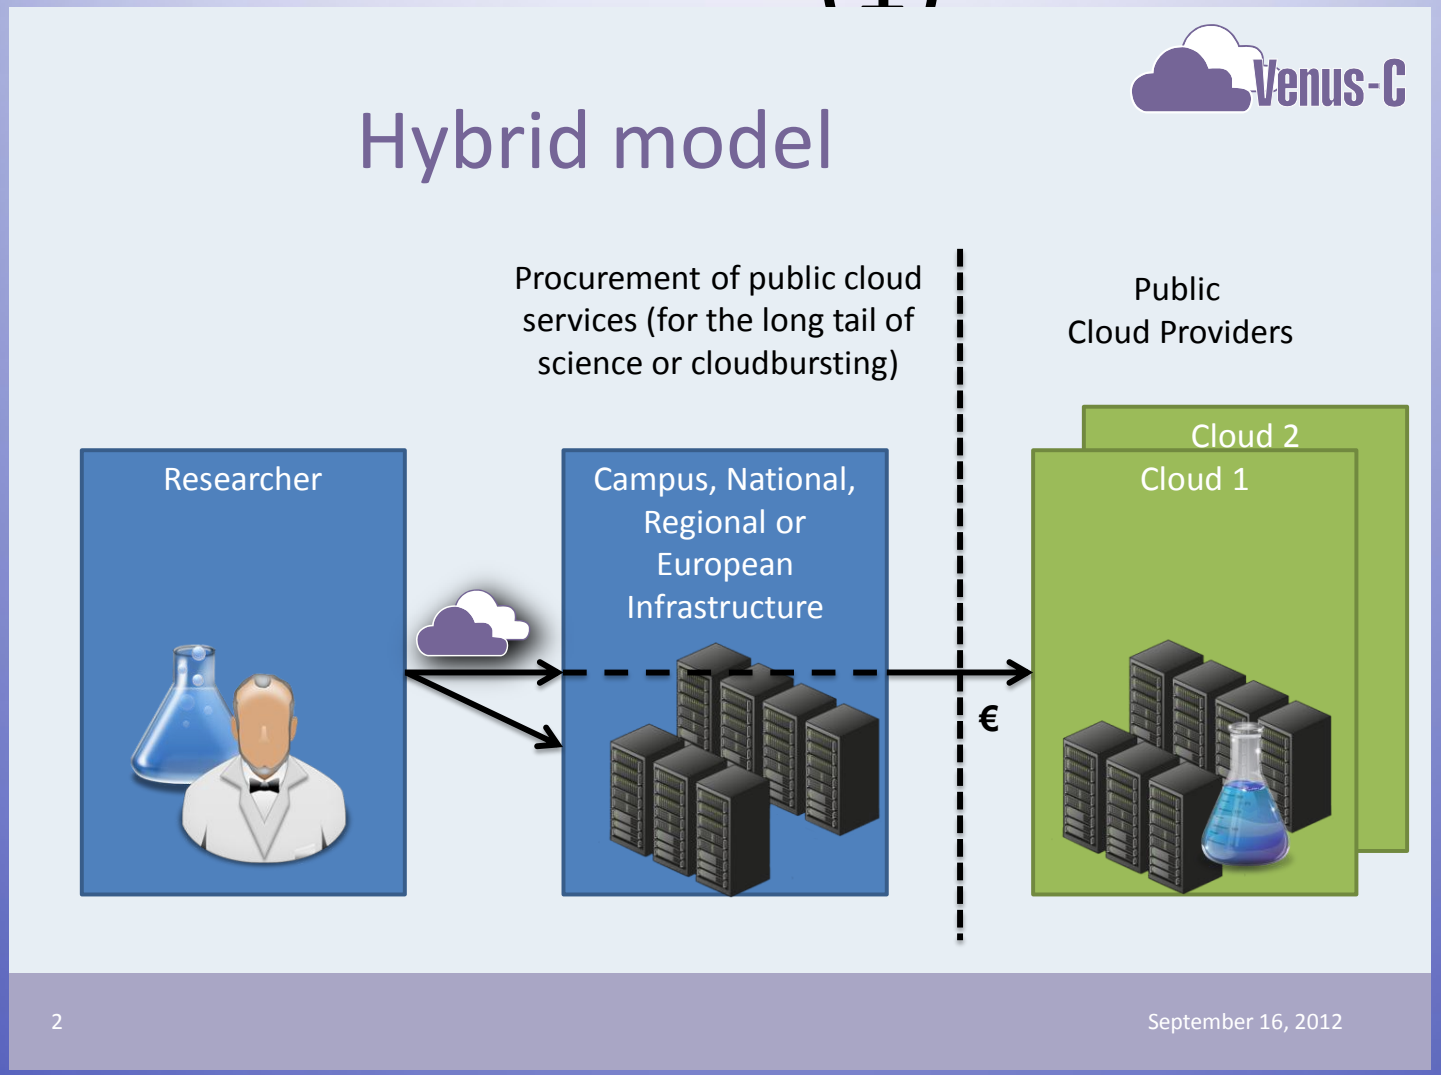

Business models for Research (2)

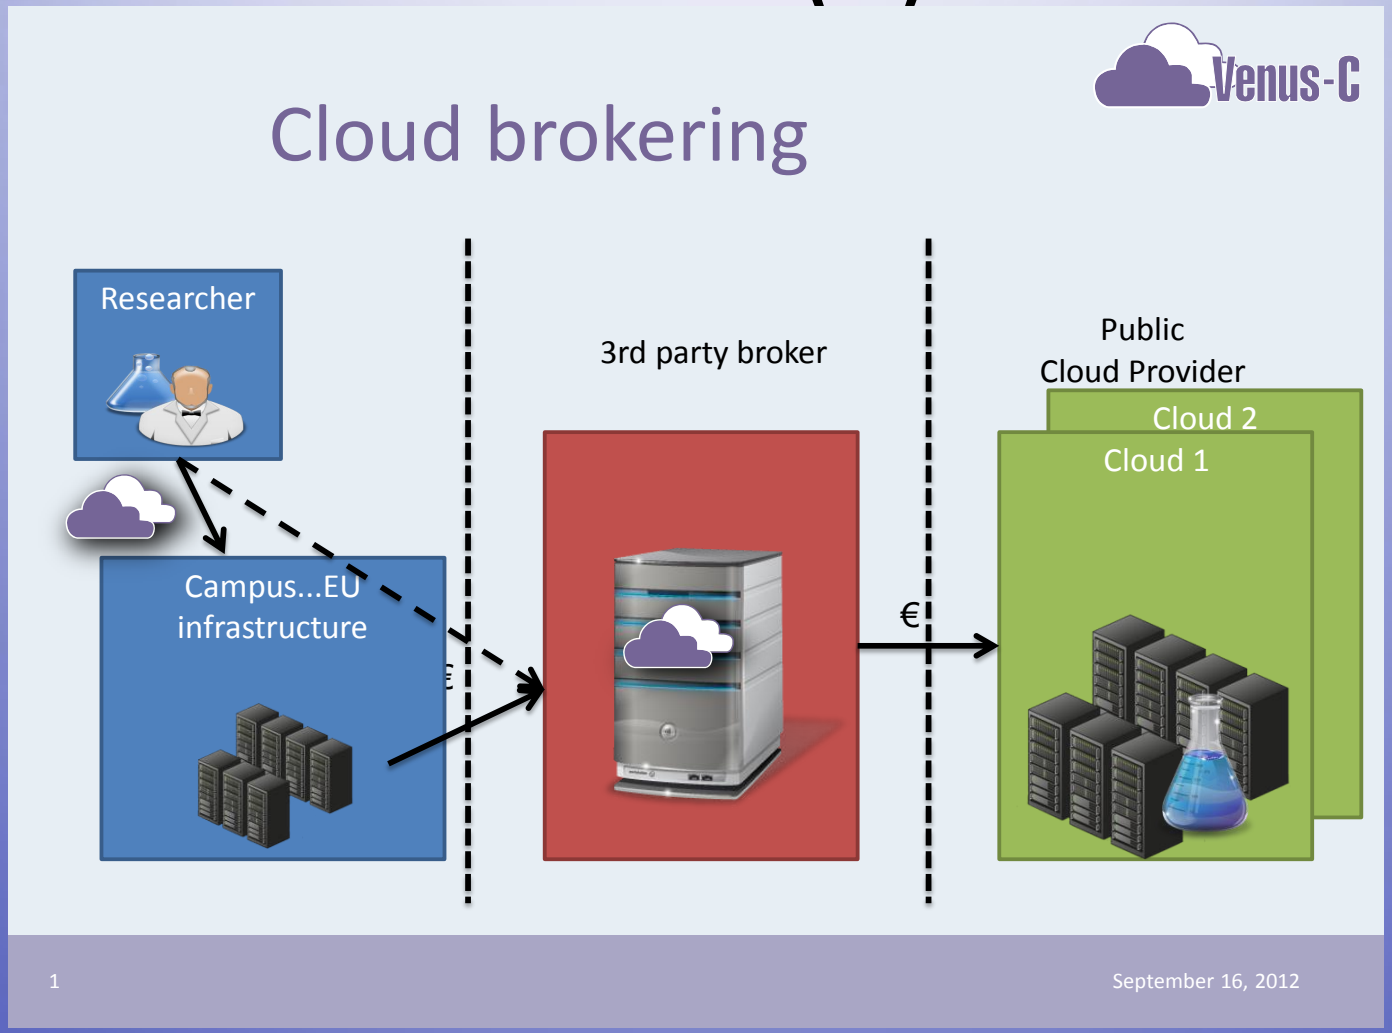

Business models for Research (3)

Individual scientists directly access public clouds