Strategy of the Commons:

Modelling the Annual Cost of Successful ICT Services for European Research

Matti Heikkurinen

matti@emergence-tech.com

e-FISCAL project/Emergence Tech Ltd.

Quick recap

- e-Infrastructures are services *based* on distributed and federated components
  - Networking
  - **Computing**
  - Storage
  - …
- Aiming at enabling excellent research
  - From Higgs to Human genome, Atoms to Arts
  - Major, predictable resource requirements
- Pervasive, pan-European
  - Equal access to all researchers

Contents

- Why assess the costs?
  - Takes 1 – 3
- Approaches and tradeoffs
  - The e-FISCAL model
- Initial results
  - Numbers – and Cloud comparison
  - Benchmarking
  - Other observations: questioning the numbers.
- Future plans and conclusions

Why assess the cost?

Why assess the costs 1/3

- "Feet on the ground" reasoning:
  - Growing use
  - New users
    - From few data producers to thousands/discipline
  - New technologies and opportunities
    - Cloud
  - Need to make more informed choices
    - To ensure that the taxpayers get the best ROI

Why assess the costs 2/3

- "Transfer argument":
  - Costs of the e-Infrastructure: "accounting grand challenge"
    - Model developed here is robust and lightweight
  - Federated ICT environments are becoming feasible
    - Also outside "big science" – e.g. Networked Enterprises WS
  - Make cost assessment more feasible for the "middle"
    - Big companies have experts
    - Start-ups have passion and hunger – no accounting in mid-leap
    - Middle-size: the awkward level
      - Know they should optimise (they see the bills)
      - High opportunity cost for optimisation: new client vs. savings
      - Can't invest the 0.5 FTE for months needed for traditional approach

Why assess the costs 3/3

- The Zen of cost assessment
  - It is the process, not the result!
  - Assessment as a trigger for broader self analysis
    - Cost of one helpdesk ticket == 1000 core hours
    - FAQs, updated manuals, peer-support save money
    - Happy users -> growing use -> happy funding agencies
  - Polar opposite of austerity measure
    - Don’t fire your community liaison to save 0.001€/core hour!

Approaches and trade-offs

Approaches: Full cost

- Full cost accounting for annual cost
  - Keep tally of the expenditure over lifetime of the “thing”
    - From submitting the application for building permit of the computing centre
    - To lighting the fuse under the building (and settling of the dust)
  - Check the calendar
    - Divide total cost by years passed
  - Main issue: latency
    - Accurate obsolete information!

Approaches: TCO

- Annual cost based on TCO
  - While writing the application for building permit reflect on
    - Lifetimes of different components
    - Maintenance requirements
    - Staffing costs decades into the future
    - Rate of obsolescence of technologies
    - Building and decommissioning costs
  - Divide the total with the longest lifetime
  - Educated guess rapidly

The e-FISCAL approach: quick measure + short projection

- Collect information about existing infrastructure
  - "What is out there?"
- Make estimate of the current replacement cost of each component
  - "What would it cost if we built it today?"
- Divide the cost by useful lifetime
  - Add OPEX

The e-FISCAL approach

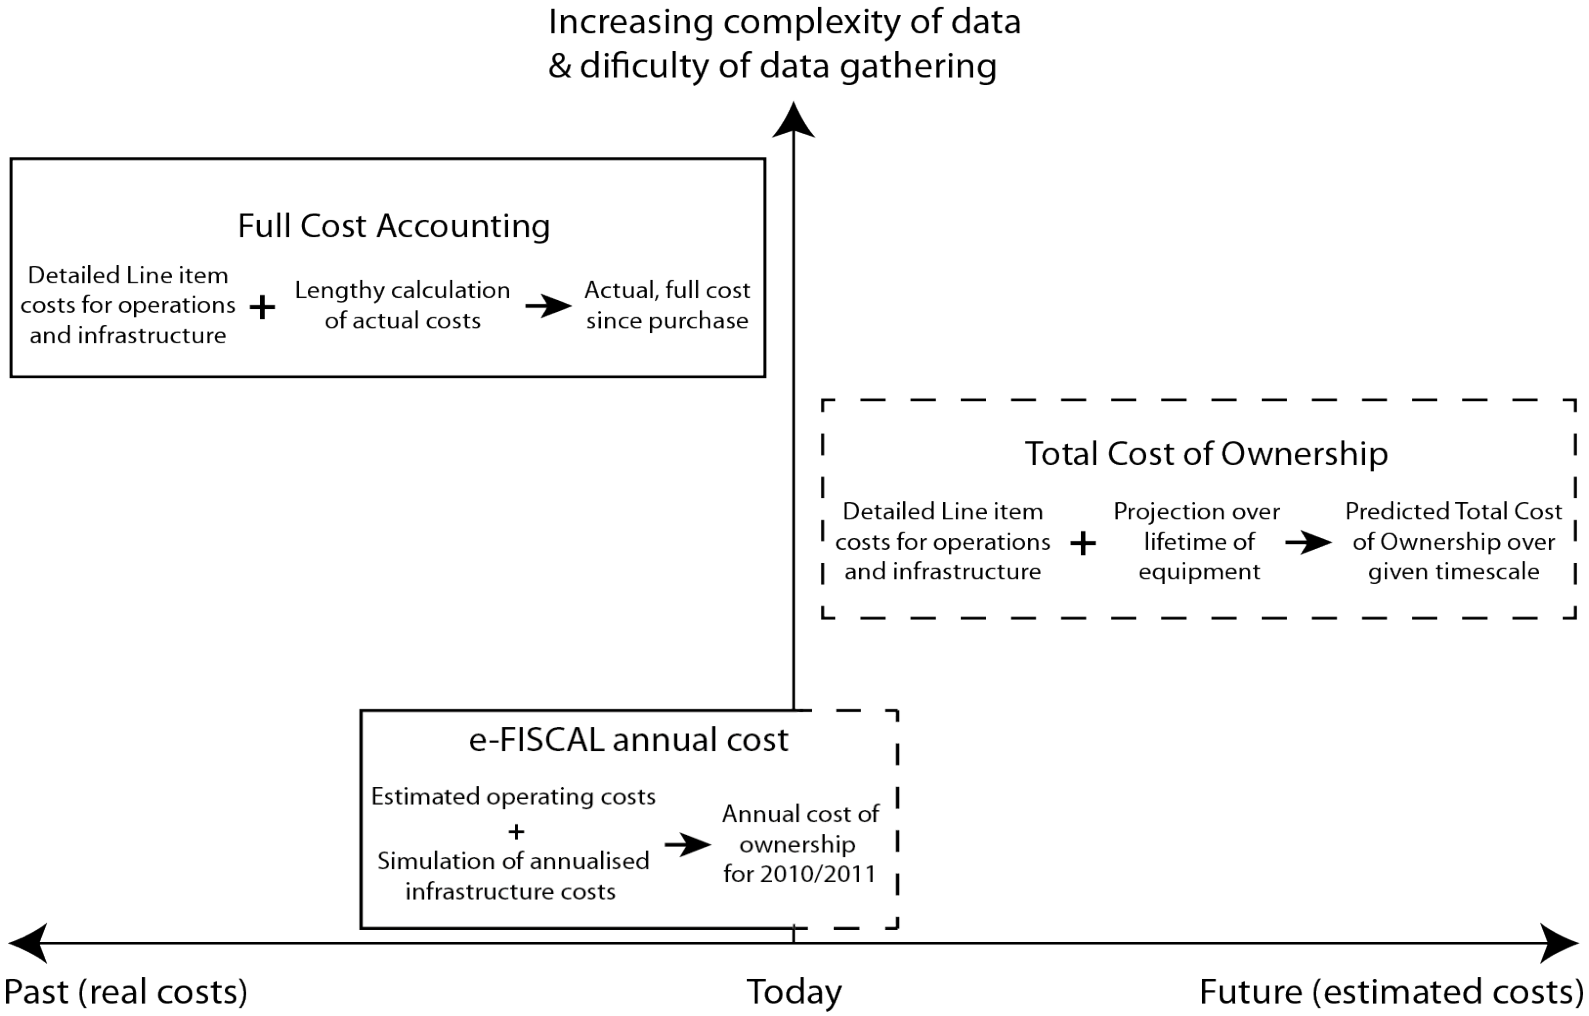

Initial results

Sample/Respondents

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

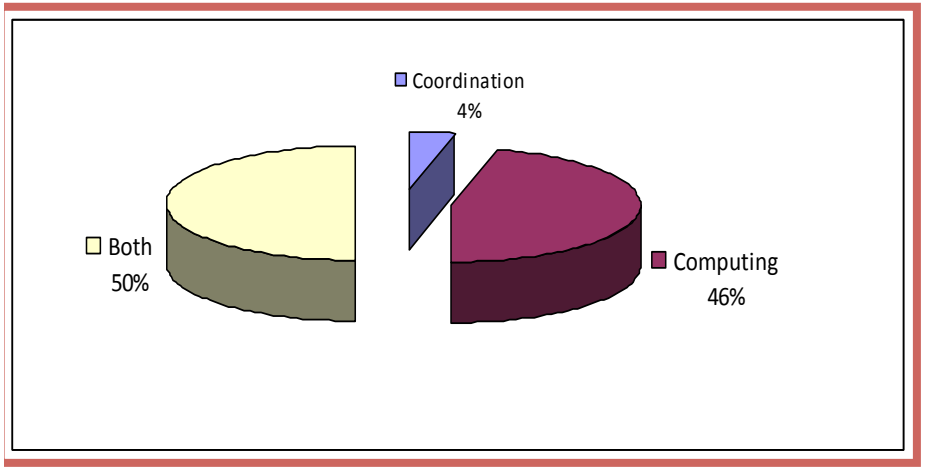

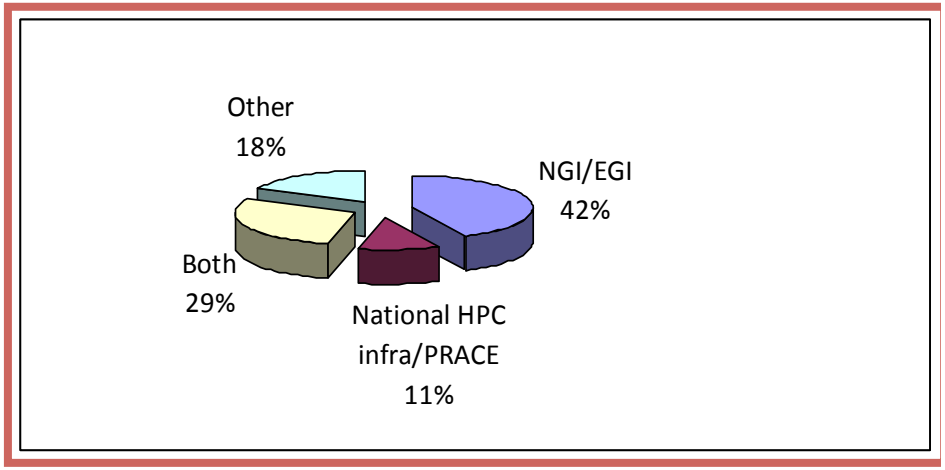

  - Majority provide both computing and coordination
  - Most of the data from HTC or mixed HTC/HPC centres

Numbers 1/2

	2010		2011	
Amounts in €	Average	Median	Average	Median
Total yearly CAPEX/ Logical CPU	119.3	86,7	104.7	61.5
Total yearly operating costs (OPEX)/ Logical CPU	396.0	208,3	290.5	140.3
Total yearly cost/ Logical CPU	515.3	295	395.2	201.7
Operating costs / total yearly costs	76.85%	70.62%	73.51%	69.54%
Capital costs / total yearly costs	23.15%	29.38%	26.49%	30.46%

Median 2011

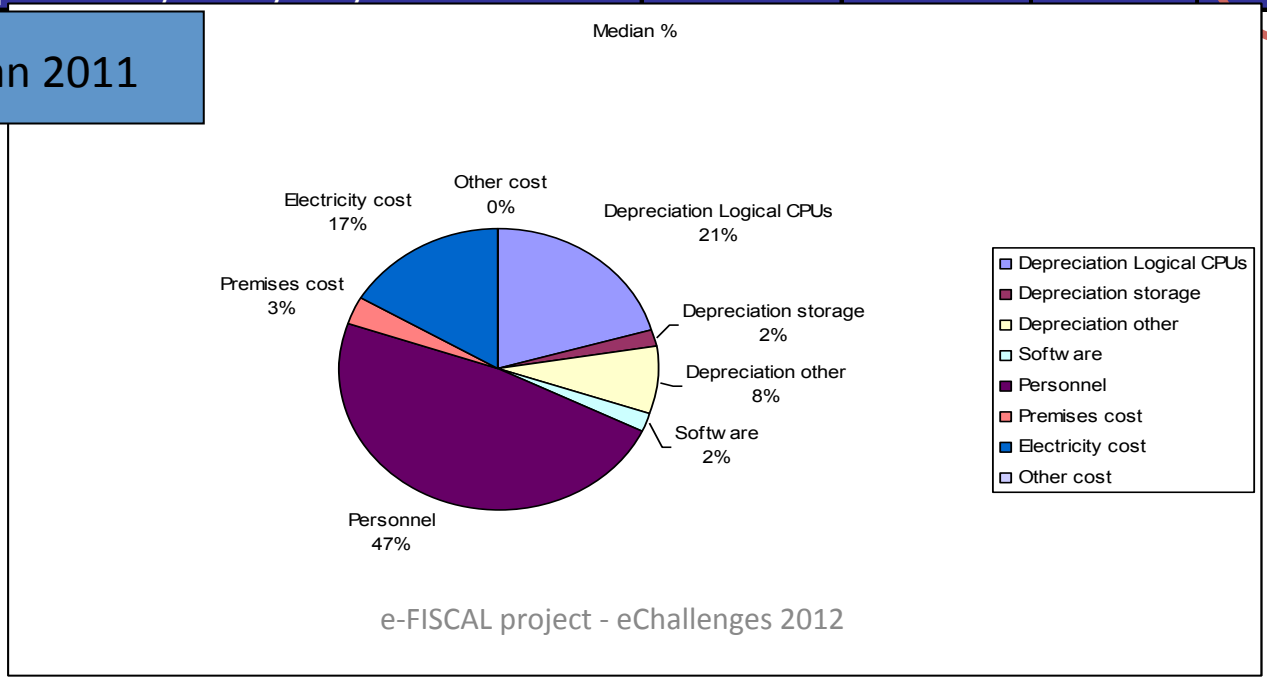

e-FISCAL vs. “orthodox cloud”

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

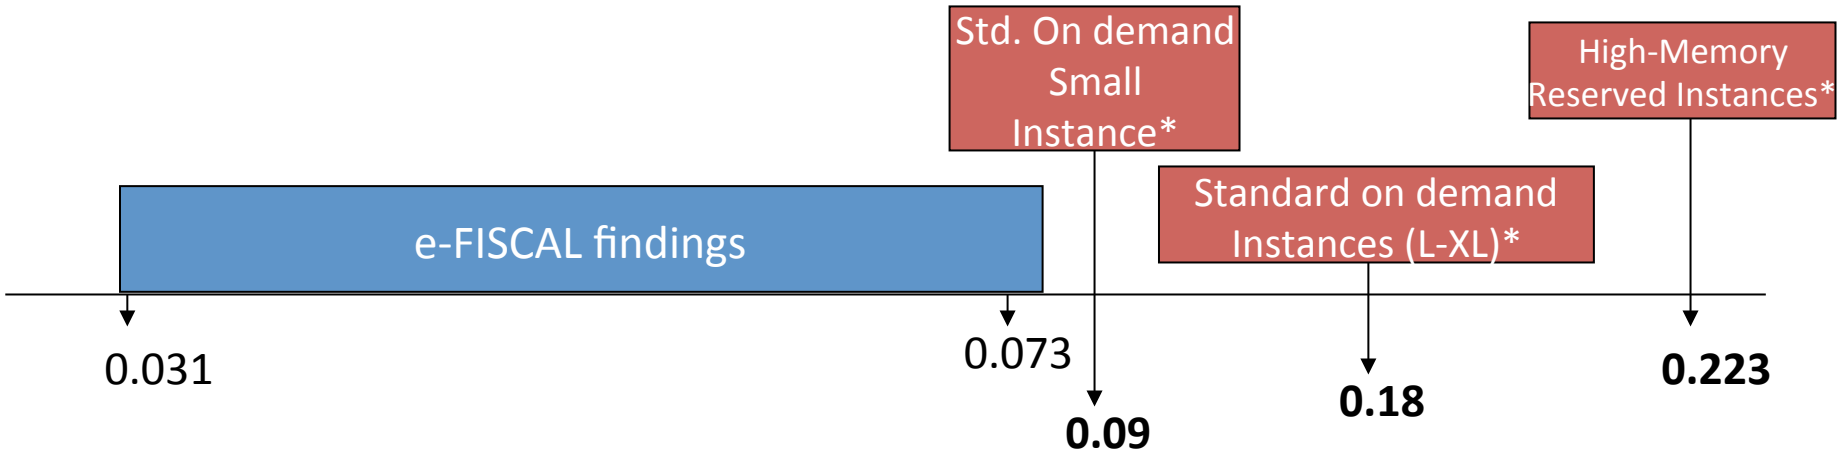

*Cost for instances/hour
transformed in €/logical CPU hour (equivalence based on instance characteristics)
Based on windows/EU-Ireland
Amazon site accessed on 12/9/2012, 1 € = $ 1,2878

Notes: **a. No performance adjustment**
b. Networking costs have been excluded in both cases
c. e-FISCAL findings include some storage costs

e-FISCAL with “Cloud leasing”

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

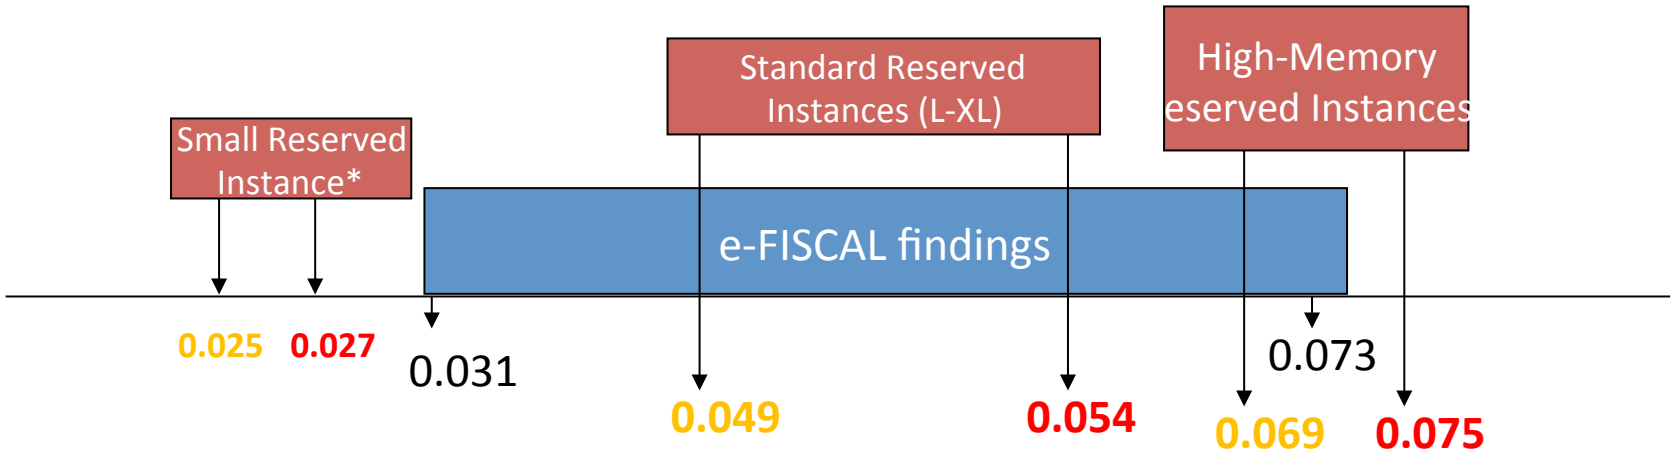

*Cost for 3-year reserved instances/hour
transformed in €/logical CPU hour (equivalence based on instance characteristics)
Based on windows/EU-Ireland/80% (red) -100% (yellow) usage of reserved instances.
Amazon site accessed on 12/9/2012, 1 € = $ 1,2878

Notes: **a. No performance adjustment has been performed YET**
b. Networking costs have been excluded in both cases
c. e-FISCAL findings include some storage costs

HS06 for Medium (SBridge)

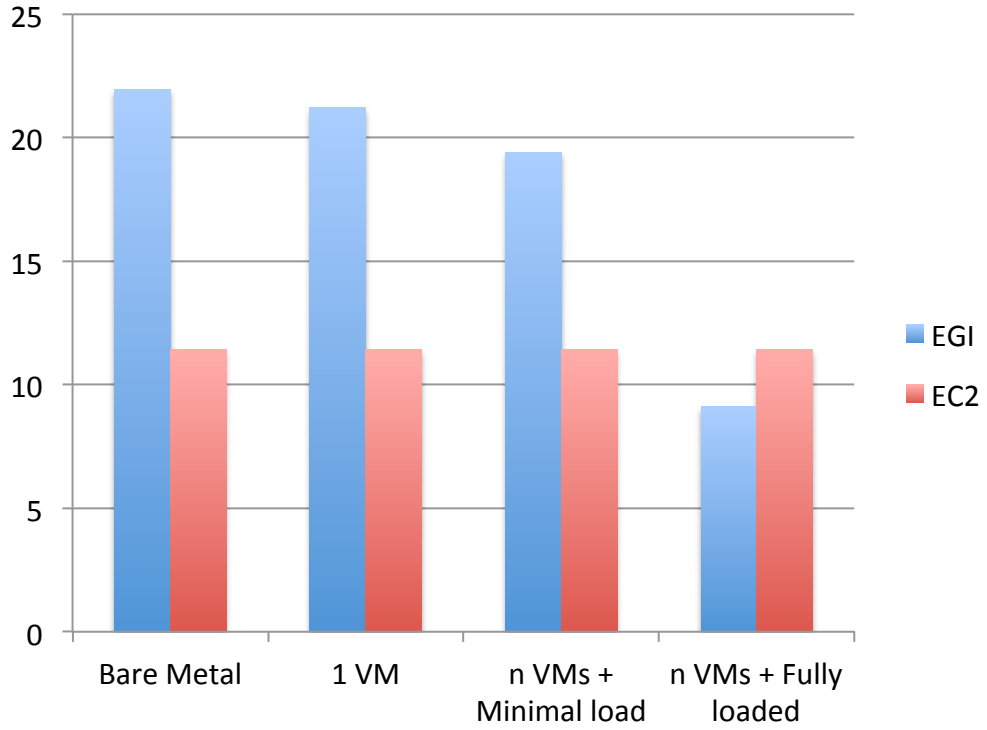

- SPEC score < with the > no. of VMs
- Virtualisation + MT effect on performance ~ 3.28% to 58.48%
- More realistic figure ~ 11.53 to 58.48

Check six!

- "e-Infrastructures are services ***based*** on distributed and federated components"
  - Clouds are services providing these components
  - Direct Cloud comparison is not Apples vs. Oranges, it is **Pines vs. Pineapples!**
- Excellent science pushes things to the limit and beyond
  - Moving from quantitative limits to qualitative
  - Qualitative limits often discovered when things break down
- e-Infrastructure as a layer to share "learning experiences"
  - Using a power switch or a credit card to switch CPUs on a minor detail
  - Literature: use of Cloud does not reduce manpower requirements from the system administration

Role of numbers

- Numbers matter
  - Basis for budgeting
  - Identifying anomalies (optimisation opportunities)
  - Measuring progress in the same centre
  - Engaging with the users (when migrating to pay-per-use model)
- They don't matter
  - Comparison with other centres
  - Demonstration of value provided
- Are awkward when
  - Comparing cost and prices
  - Attempting to counter hype

Future plans and Conclusions

Future plans

- The project will continue until early 2013
  - Final workshop
- Wrap up
  - Conclude survey and analysis
  - Polish the reusable cost models and tools
- Handover
  - EGI compendium effort
  - Online community
  - Other interested parties?

Conclusions

- Cost assessment in federated environments
  - Challenge and opportunity
- The e-FISCAL approach is not tied to the e-Infrastructures
  - Excellent pilot environment
- Cost discussions may be hard
  - Results of not having them considerably harder!

e-FISCAL

Thank you for your attention!
Questions?

For more information:

info@efiscal.eu - www.efiscal.eu

Send an e-mail to join the mailing list and to hear about the final workshop!

Co-financed by the EC under Grant Agreement Number 283449

Cost structure and yearly cost

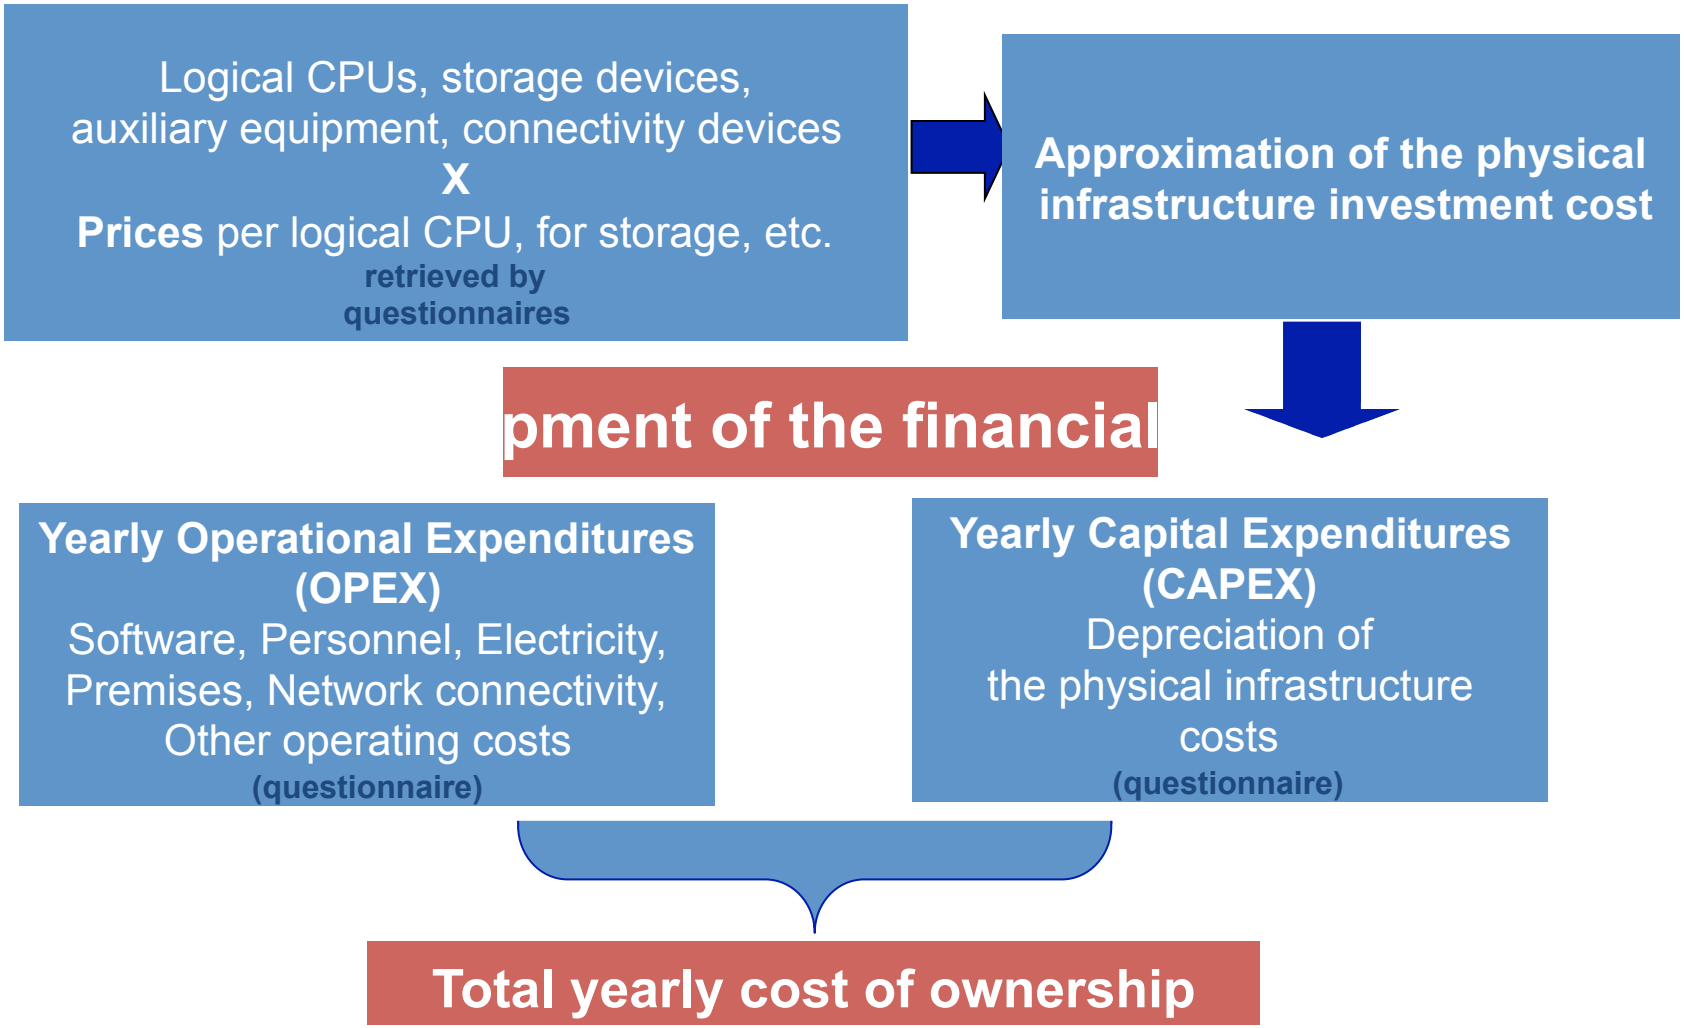

Summary numbers

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)	€ 0.075	Study for JISC UK - Differences between institutions reviewed
US DoE - Magellan report (2011)	€ 0.015	Hopper system – National Energy Research Scientific Computing Centre- including storage sub- system
Smith (2011)	€ 0.031	Purdue campus, USA
University of Washington	€ 0.020	Hyak cluster, USA
Cohen and Karagiannis (2011)	€ 0.09 – € 0.14	Stratified sample of EGI centres - Assuming 60% utilization ratio – storage cost included ((numbers refer to 2009)
Cohen and Karagiannis (2011)	€ 0.08 – € 0.10	Stratified sample of EGI centres - Assuming 60% utilization ratio – storage cost excluded (numbers refer to 2009)