

Assessing the cost of HPC and HTC infrastructures in Europe: The e-FISCAL cost findings

Sandra COHEN, Fotis KARAGIANNIS¹, Matti HEIKKURINEN², Kashif IQBAL³ Sergio ANDREOZZI, Sy HOLSINGER⁴

¹ Athens University of Economics and Business, Greece, *Email: {scohen, fkara}@aueb.gr*

² Emergence Tech Ltd, UK, *Email: info@emergence-tech.com*

³ Irish Centre for High End Computing (ICHEC), Ireland, *Email: kashif.iqbal@ichec.ie*

⁴ EGI.eu, The Netherlands *Email: {sergio.andreozzi, sy.holsinger}@egi.eu>*

Over the last ten years the European countries and the EC have made significant investments in e-Infrastructures for scientific computing, notably High Throughput Computing (HTC) and High Performance Computing (HPC) services. Their contribution in pushing the European research towards a leading position in addressing global challenges has been validated by a growing number of research initiatives. Sustainability of such services is essential, as the research supported by them is more and more crucial for European competitiveness – and sustainability can be planned better if the costs are known. The emerging commercial offerings (Cloud-based HTC and HPC solutions) pose additional challenges and opportunities for sustainability. It is thus important to understand the cost of the dedicated computing related e-Infrastructures, as expressed by EGI and PRACE initiatives. The calculation of such costs is not trivial; keeping detailed accounting data for the evolving mix of capital (hardware) and operational (personnel, energy) expenditures is only one of the challenges.

The goal of the e-FISCAL initiative is to analyse the costs of the current European dedicated HTC and HPC computing e-Infrastructures for research and compare them with equivalent commercial leased or on-demand offerings.

In doing so, e-FISCAL develops a hybrid costing methodology that builds on the two main methodologies used for cost assessment in e-infrastructure: Total Cost of Ownership (TCO) and Full Cost Accounting (FCA). Moreover, it capitalises on the state-of-the-art literature. While TCO is a useful concept in assessing the cost of a specific infrastructure or project over its useful life it is highly demanding in analytical data that have a forward looking orientation in order predictions to be acceptably precise. On the other hand, FCA methodology relies on actual cost accounting data information that is dependent on the level of cost accounting systems sophistication by additionally inducing a backward looking stance.

e-FISCAL model is a hybrid model that approximates the costs of maintaining services at their current level in the short to medium term. It does that without the need to identify funding sources or the exact points in time when the actual infrastructure investments have been made. Therefore, rather than relying on the detailed financial data or projections to the future, it uses high level information about computing and storage hardware costs (including interconnection costs), auxiliary equipment costs (i.e. cooling, UPSs, power generators), software costs, personnel costs, and site operating costs. The e-FISCAL methodology is completed in two phases; firstly a simulation of the physical infrastructure is sketched and secondly the annualised cost of the simulated physical infrastructure and the operating cost of the physical infrastructure are added together. All the data necessarily to feed the model (e.g. e-infrastructure acquisition costs, personnel

costs, electricity cost, depreciation rates) is retrieved through a properly developed questionnaire (the e-FISCAL survey instrument).

The e-FISCAL methodology has been applied to a sample of HTC and HTC centres¹ in Europe that contributed to the study, making therefore, the e-FISCAL study the first to gather and analyse the costs from 14 countries in a comprehensive and systematic way. As it is evident by the state-of-the-art review, the majority of costing studies either concentrate on a single site or on multiple sites yet in the same country. The analysis revealed that the median² cost per core/hour for 2011 is € 0.03/core hour while the corresponding average value goes up to € 0.07 /core hour. Some other underlying cost dimensions that govern the results above and were revealed by the analysis of input data are that: a) Hardware depreciation rates are in several cases well above the typical well-cited three –year period. The average depreciation period corresponds to 5 years. b) Operational expenditures (OpEx) dominate the annual costs of the sites (69%) over the depreciated capital expenditures (CapEx) per year (31%) c) Personnel costs constitute almost 50% of the overall annual costs. d) The Power Usage Effectiveness (PUE) rates are around 1.5 (median value) indicating rather efficient energy use and e) Approximately 2.12 FTEs correspond to 1.000 cores (median value). Moreover, the analysis provided evidence of a decreasing trend in costs from year 2010 to 2011.

After assessing the costs per core/hour a comparison between the prices charged by commercial cloud providers (e.g. Amazon EC2) and costs estimated through e-FISCAL has been conducted. This comparison has been performed two-fold; firstly without taking into consideration any performance differences between in-house service provision and cloud and secondly, after adjusting for differences in performance (performance-adjusted cost comparisons). In order to execute the second type of analysis, we rely on a short-scale benchmarking exercise that was conducted for the scope of the project. The latter demonstrated an average 40% performance degradation of the Amazon instances compared to comparable HPC and HTC site instances.

By consternating only on financial considerations, our findings provide corroborative evidence that cloud prices are not necessarily and unanimously lower than the costs calculated by e-FISCAL of in-house service providing. Parameters like utilization rates, infrastructure size, length of time commitment in rendering cloud services, efficiency in manpower utilization and service performance are important factors affecting the cost relation between cloud and e-infrastructures. Our study also touches upon non-financial aspects that are relevant to the cloud vs. in-house discussion such as the speed to adjust to increases in capacity demands, the easiness and willingness of porting applications from one environment to another, security considerations as well as the facilitation of knowledge transfer and expertise related to leading edge ICT challenges. This latter discussion acknowledges that the cost is not synonymous to value; on the contrary cost is only one parameter in the value creation process.

¹ Our analysis is based on 26 answers. Twelve respondents indicated that their institute is only part of NGI/EGI, 3 reported participating into National HPC infrastructure/PRACE, while another 10 participate in both.

² The median is described as the numeric value separating the higher half of a sample, a population, or a probability distribution, from the lower half. Median is more efficient in a sample with a wide range of distributions reducing the impact of extreme low and high values.

Related presentation: [http://www.efiscal.eu/files/presentations/Karayannis - e-FISCAL Workshop @EGI TF - September 2012 - Prague-vfinal.pdf](http://www.efiscal.eu/files/presentations/Karayannis_-_e-FISCAL_Workshop_@EGI_TF_-_September_2012_-_Prague-vfinal.pdf)