

HPC Opportunities and Challenges in e-Science

Fabrizio Gagliardi

Microsoft Research, Technical Computing,
EMEA and LATAM Director
Phone: +41 78 844 6476
fabrig@microsoft.com

Abstract. In the last few years a new paradigm has emerged in science: extensively use of simulation techniques and software modeling, running on a distributed high performance computing electronic infrastructure. This paradigm is referred to as electronic Science or e-Science. Besides computer simulation, it uses huge amounts of distributed and shared data captured by instruments or sensors and/or stored in databases, analyzed to provide new results for science. This distributed HPC and data ecosystem enables the sharing of acquired knowledge, remote access to resources and above of all a world-wide scientific collaboration.

Keywords: e-Science, Grid computing, Cloud, elastic computing, multi-core architectures.

The EU Datagrid project was one of the first examples of Grid computing going in this direction, which was followed by the EGEE (Enabling Grids for E-science) project. The research community and the High Energy Physicists at CERN were among the first to adopt Grid computing. The EGEE projects has to date integrated more than 50.000 CPUs in Europe and beyond, and 20 Peta Bytes (millions of GigaBytes) of storage, serving multiple application communities including HEP, Bioinformatics, Astrophysics, Computational Chemistry, Earth Sciences, Fusion. Some business/industrial applications are also adopting this distributed HPC computing model such as automotive, finance, multimedia, and there also a few examples of e-Government ones such as in the civil protection area.

Thus, Grid computing has delivered an affordable and high performance computing infrastructure to scientists all over the world for solving intense computing problems within constrained research budget. Business or industrial entities have also used similar technologies to increase the usage of their computing infrastructure and reduce their total cost of ownership (TCO). In addition, Grid computing is leveraging the advanced research networks to deliver an effective and irreplaceable channel for international collaboration.

Issues which hinder a wider adoption of grid technology in e-Science and industry have to do with the cost of operations and the overall complexity of the Grid, which aims at delivering secure and reliable services over widely distributed non homogeneous resources belonging to different administrative domains. The EGEE project, for instance, is spending more than 30 million € per

year (around half is covered by the EC) to run and support the 50.000 CPUs infrastructure (operations, middleware development and certification, application support, training, dissemination, etc.). Besides the operational expenditures, one should take into account the depreciated capital expenditures of the infrastructure. Hardware costs of the 50.000 CPUs and the 20PB storage should be in the order of 100 million Euros, depreciated over 5 years, results for another 20 million Euros per year.

Power consumption and heat dissipation are also becoming an important new factor that needs to be considered seriously. Taking the rough estimation of 10% of the hardware costs, this should add another 10 Euros per year operational expenditures. Supposing that the (over-provisioned) connectivity costs are covered anyway by the National Research Networks, this sums up to around 60 million € per year. This is comparable to the estimated 48M € figure that the EGEE usage would have cost, if performed with Amazon Elastic Computing and Simple Storage Services (presented by the EGEE project director in the 2008 EGEE User Forum). Although calculations are not accurate, they show that the grid and the cloud now are in the same order of costs.

Thus elastic computing, Computing on the Cloud, Data Centres and Service hosting are offering on demand CPUs and storage with similar pricing. As an example, Amazon Elastic Computing (EC2) offers a small instance CPU hour for 0,10 USD, and is quite easy to use. In addition, it charges 0,10 USD per GB for data transfer in and 0,18 USD for data transfer out (of their systems). For datasets that stay at the Amazon system, one has to use the Amazon S3 services, which have an additional cost of 0,15 USD per GB per month. Other major stakeholders in the market such as SUN, Google and IBM are moving in the same direction towards offering similar services of on-demand computing, storage and hosting.

Emerging multi-core architectures and CPU accelerators promise potential breakthroughs, and in the future one might not have to rely on computer clusters, the cloud and the grid, rather afford her/his own personal supercomputer desk top or desk side

Microsoft is actively investigating this field, and the Technical Computing in Microsoft Research, is supporting e-Science initiatives in collaboration with leading scientists around the world. We need to advance in making computing easy to use for the scientists to concentrate their energy in real science and not the computing tools!