Toward a Flexible, Environmentally Conscious, On Demand High Performance Computing Service

G.B. Barone, R. Bifulco, V. Boccia,
D. Bottalico, R. Canonico
University of Naples Federico II
Italy
gbbarone@unina.it, roberto.bifulco2@unina.it,
vania.boccia@unina.it, davide.bottalico@unina.it,
roberto.canonico@unina.it

L. Carracciuolo
Italian National Research Council
Italy
luisa.carracciuolo@cnr.it

***Abstract*—This work is related with planning and implementation of an on demand computing service which is able to obtain a right trade-off between managment cost reduction, environmental sustainability and user satisfaction.**

***Keywords*-internet services; computing on demand; green IT; cloud computing; virtualization; High Performance Computing;**

I. Introduction

On demand computing is a model in which computing resources are made available to the user as needed. It could be considered a valid solution for people who need a huge amount of resources, to reduce the *Total Time to Solution*, and cannot bear the costs of HPC systems. Those costs are related i.e., to energy consumption, cooling systems and specialized know how for hardware/software resources upgrade and maintenance.

On the other hand, from the point of view of resources manager, it is important to find the right trade-off between an effective use of resources and the cost reduction also in terms of energy consumption.

Certainly, both users and resources manager, would benefit by environmental sustainability of new *environment-conscious* HPC systems.

The purpose of this paper is to describe our first experience in designing and implementing a flexible infrastructure, built on the basis of both physical and virtual resources, in the name of saving energy, and overall cost reduction, and than in the name of a more efficient resources usage [7]. We note that the term "flexible" is used to indicate an infrastructure where the number and type of resources may change on the basis of the user requirements and real workflow.

In section II we describe the architecture of planned on demand computing service.

In section III we describe a Case Study related with the implementation of the service on the SCoPE Computing Infrastructure of University of Naples[6].

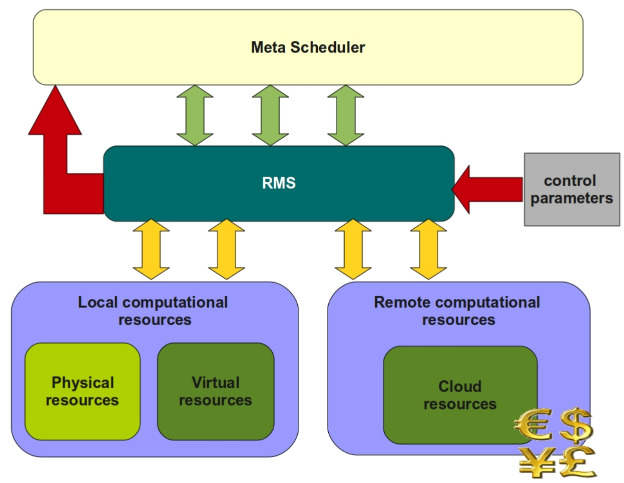

Figure 1. Service schema

In section IV we give some information on our future work.

II. Service planning

Local or distributed large scale systems offer on demand computational services to different, and often heterogeneous, user communities. These services usually have to meet constraints, defined in SLAs (Service Level Agreements), that need resources are always *"online"* even if they aren't effectively used.

However, computational resources power supply and cooling affect very heavily on total energy consumption and then on management costs. Therefore large scale systems owners have to find out a compromise between the sustainability of such systems and overall user community satisfaction.

We just worked on planning and deployment of a software solution, whose architecture is represented in figure 1, that implements an energy-aware Resource Management

System (RMS), able to dinamically re-define the set of physical/virtual *”online“* computational resources.

In figure 1, Metascheduler recieves inputs from user community and, on the basis of SLAs, notify the RMS with needed resources allocation requests.

The RMS works on the bases both of these requests and of some control parameters. Control parameters represent a solution to an optimization problems: maximizing or minimizing some metrics functions (i.e., absolute energy consumption, time to solution/energy consumption, etc.), with the aim to setup the most affordable set of computational resources. The above described process produces a suitable combination of physical, virtual or cloud resources, that is finally notified to the Metascheduler.

We notice that the selection of physical/virtual local or cloud resources may depends also on their actual *”price“*: for a local resource we consider the total energy consumption, while for a cloud one we intend the purchase cost of a green awareness IT resource.

III. Service deployment: a test case

The reference infrastructure is S.Co.P.E. Datatcenter at University of Naples Federico II. It consts of up to two thousands computational cores.which represent about 90% of the datacenter hardware, therefore they affect very heavily on management costs in terms of energy consumption for cooling and power supply.

S.Co.P.E. resources are integrated into local, national (GRISU‘ and IGI grids) and international (EGI grid) relevant distributed computational infrastructures, on the bases of memorandum of understanding (MOUs). MOUs contain agreements about the service level to be granted (SLAs): availability and reliability minimum levels for network, computational and storage resources, the amount of cores and disk space to be shared, etc.

All computational resources are always powered on even if not used, so our work is strategic to ensure the sustainability of the above infrastructure.

In figure 2 we show how the computing service has been implemented on S.Co.P.E. infrastrucure.

From bottom up, firstly we integrated into reference infrastructure some virtual and cloud resources: virtual resources are provided by means of XEN [4] virtualization technology and related tools, while cloud resources are based on Eucalyptus [5] system which is an Amazon EC2 compatible cloud middleware.

We then worked on the integration of the local RMS (SCoPE RMS) with a decision engine and a set of protocol interfaces.

Decision engine select the most suitable combination of physical, virtual or cloud resources to be *“powered on”* on the bases of workloads submitted to Scheduler, defined SLAs and values, or estimates, of some metrics that measure *power-perfomance efficiency* of application(i.e. Energy-Delay-Product metric [3]) and/or the total sytem energy consumption[2].

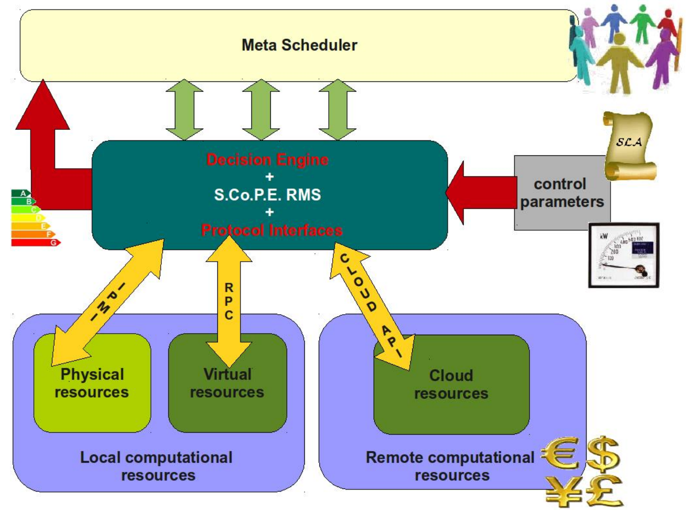

Figure 2. Service deployment on the S.Co.P.E. Infrastructure

Resources are *“powered on”* by protocols interfaces becoming available to the Metascheduler. Depending on the type of resources, three type of interfaces are used: IPMI (Intelligent Platform Management Interface) for physical resources, RPC (Remote Procedure Call) for local virtual resources and Eucalyptus API/tools for cloud resources.

IV. Conclusion

In this paper we describe our first experience in designing and implementing a flexible infrastructure, built on the basis of both physical and virtual resources, in the name of saving energy and than in the name of a more efficient resources usage. Our future work will be also related with the definition of some *best practice* in building flexible HPC infrastructure which could be a good trade-off between managment cost reduction, the environmental sustainability and a new way of understanding the user satisfaction [1].

References

[1] I. Foster, *Whats faster-a supercomputer or EC2?* [1], 2009.

[2] S.K. Garg and C.S. Yeo and A. Anandasivam and R. Buyya, *Energy-Efficient Scheduling of HPC Applications in Cloud Computing Environments* [2], 2009.

[3] F. Xizhou and G. Rong and K.W. Cameron, *Power and Energy Profiling of Scientific Applications on Distributed Systems*, Proceedings of 19th IEEE International Parallel and Distributed Processing Symposium, 2005.

[1] `http://ianfoster.typepad.com/blog/2009/08/whats-fastera-supercomputer-or-ec2.html`

[2] `http://arxiv.org/pdf/0909.1146v1`

[4] Barham, P. and Dragovic, B. and Fraser, K. and Hand, S. and Harris, T. and Ho, A. and Neugebauer, R. and Pratt, I. and Warfield, A. *Xen and the art of virtualization*, ACM SIGOPS Operating Systems Review, volume 37, number 5, 164-177, 2003.

[5] D. Nurmi and R. Wolski and C. Grzegorczyk and G. Obertelli and S. Soman and L. Youseff and D. Zagorodnov, *The Eucalyptus Open-source Cloud-computing System*, in Proceedings of 9th IEEE International Symposium on Cluster Computing and the Grid, 2009.

[6] L. Merola, *The S.Co.P.E. Project*, in Proceedings of the Final Workshop of Grid Projects "PON RICERCA 2000-2006, AVVISO 1575".

[7] R. Talaber and T. Brey and L. Lamers, *Using Virtualization to Improve Data Center Efficiency*, The Green Grid White Paper n.19 [3], 2009.

[3] http://www.thegreengrid.org/~media/WhitePapers/White%20Paper%2019%20-%20Using%20Virtualization%20to%20Improve%20Data%20Center%20Efficiency.ashx?lang=en