Managing Cloud Resources for Medical Applications

Piotr Nowakowski[2], Tomasz Bartyński[2], Tomasz Gubała[2], Daniel Harężlak[2], Marek Kasztelnik[2], Jan Meizner[2], Marian Bubak[1,2]

[1]AGH University of Science and Technology, Department of Computer Science, al. Mickiewicza 30, 30-059 Kraków, Poland
[2]AGH University of Science and Technology, ACC Cyfronet AGH, Nawojki 11, 30-950, Kraków, Poland

emails: bubak@agh.edu.pl, {p.nowakowski,t.bartynski,d.harezlak,t.gubala,m.kasztelnik,j.meizner}@cyfronet.pl

Keywords: cloud infrastructures, eScience, medical applications

1. Introduction

The synergy between computational clouds and non-HPC scientific applications (sometimes referred to as "midrange computing") is highlighted in recent institutional reports concerning, such as the Magellan Report on Cloud Computing for Science [1] and the CERN Strategic Plan for a Scientific Cloud Computing Infrastructure for Europe [2]. From a scientific standpoint, the goal of the presented work is therefore to address issues which hamper the exploitation of clouds in support of medical research within the VPH community. The objectives include detailed analysis of medical application requirements with regard to data security and curation, designing a suitable platform, enabling application services to be shared in a controlled way and performing assessment of the system.

2. Description of the platform architecture

The approach adopted in the design and implementation of the VPH cloud computing platform bases on the authors' prior work on cloud resource management [3]. The cloud platform covers the entire application development lifecycle, from inception to scalable exploitation, assisting each participating class of user at each step. The VPH cloud platform must form a bridge between the world of cloud middleware services and the familiar OS environments in which standalone scientific applications are deployed. It acknowledges the difference between public and private clouds. For instance, patient Electronic Health Records are usually not permitted to leave their originating institution and must therefore be maintained on resources which comprise a local cloud site. The cloud platform must be aware of these restrictions and respect them when deciding upon service deployment.

Atmosphere, the architecture of which is depicted in Fig. 1, provides the ability to tag selected Atomic Services as "static" which means that at least one instance of a service is kept online at all times, ready to serve user requests in a rapid fashion. For other types of services Atmosphere acts as a factory: it instantiates services when requested by end user or by workflow development tools and returns the instance endpoint to the requestor.

Atmosphere also provides a dedicated set of user interface services for each user class. All of its UIs assume the form of web applications and can be either served as standalone tools or integrated into portals. For the purposes of the VPH collaboration, Atmosphere management portlets are aggregated by the so-called Master Interface which provides an entry point to the system for all members of the VPH community.

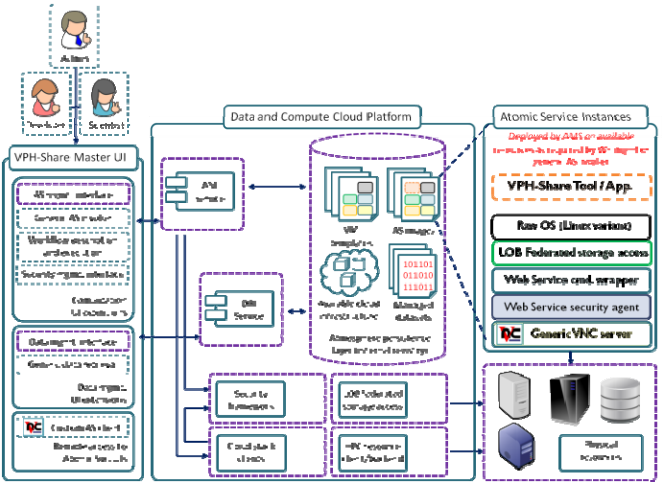

Fig.. 1. Architecture of the Atmosphere cloud platform.

3. Support for medical application workflows

At present Atmosphere is used to deploy and expose four VPH applications, with support for approximately twenty Atomic Services derived from VPH NoE applications (brain aneurism simulations, cardiovascular modeling, osteoporosis prediction tools and HIV genome analyzers). Additional applications are being deployed (including tools coming from partner projects such as P-Medicine) and the platform itself is sufficiently generic to support most e-Science workflows and standalone tools.

4. Conclusions and future work

The new PaaS solution presented in this paper is already functional, which means that it supports the full development and instantiation cycle of Atomic Services, enabling users and workflow management tools to take advantage of application components. Even so, there are extensions and new features to be considered as part of our future work on Atmosphere. Further work will focus on performance issues and support for additional cloud stack releases.

Acknowledgements. This research is partially funded by the EC ICT Project VPH-Share No 269978. The authors are indebted to Professor Rod Hose and Steven Wood of the University of Sheffield for fruitful discussions.

References

1. K. Yelick, S. Coghlan, B. Draney, and R.S. Canon, "The Magellan Report on Cloud Computing for Science", U.S. Department of Energy Office of Science, Office of Advanced Scientific Computing Research (ASCR) December, 2011; available at `http://science.energy.gov/~/media/ascr/pdf/program-documents/docs/Magellan_Final_Report.pdf` (accessed 17 July 2012).
2. Strategic Plan for a Scientific Cloud Computing infrastructure for Europe, published by CERN/ESA, CERN-OPEN-2011-036; available at `http://helixnebula.commpla.com/uploads/11/CERN-OPEN-2011-036-pdf.html` (accessed 17 July 2012).
3. M. Malawski, T. Gubała, and M. Bubak: Component-based approach for programming and running scientific applications on grids and clouds. The International Journal of High Performance Computing Applications, vol. 26, no. 3, pp. 275-295 (2012).