

# CLOUD SOLUTIONS FOR HIGH PERFORMANCE COMPUTING: OXYMORON OR REALM?

*Draško Tomić, Dario Ogrizović, Zlatan Car*

Preliminary notes

In the last years a strong interest of the HPC (High Performance Computing) community raised towards cloud computing. There are many apparent benefits of doing HPC in a cloud, the most important of them being better utilization of computational resources, efficient charge back of used resources and applications, on-demand and dynamic reallocation of computational resources between users and HPC applications, and automatic bursting of additional resources when needed. Surprisingly, the amount of HPC cloud solutions related to standard solutions is still negligible. Some of the reasons for the current situation are evident, some not so. For example, traditional HPC vendors are still trying to exploit their current investment as much as possible, thus favoring traditional way of doing HPC. The next, although virtualization techniques are developing in ever increasing rate, there are still open questions, like scaling of HPC applications in virtual environment, collaboration between physical and virtual resources, and support of the standard HPC tools for virtualized environments. At last but not the least, the advent of GPU computing has raised difficult questions even in the traditional and well developed HPC segment, like scalability, development of HPC GPU based applications. In addition, because many HPC vendors are speaking about already having a fully functional HPC cloud solution, there is a need to provide answers to the following questions: Are they fulfilling virtualization promises, both physical and virtual? Which type of clouds do they support: private, public or both, i.e. hybrid clouds? How well HPC applications scale on their cloud solutions? This paper is an overview of the current HPC cloud solutions, for sure not complete and solely a view of authors, however intended to be a helpful compass for someone trying to shift from standard HPC to large computations in cloud environments.

**Keywords:** *an overview of HPC Cloud solutions; first, second and third tier HPC Cloud solutions; future development and trends in HPC Cloud solutions*

## Super računalstvo u oblacima: oksimoron ili realitet?

Prethodno priopćenje

U posljednje je vrijeme zamjetan veliki interes HPC (High Performance Computing) zajednice prema Cloud Computing paradigmi, odnosno računalstvu u oblacima. Mnoge su prednosti izvođenja HPC aplikacija u računalnom oblaku, od kojih su najvažnije bolje korištenje računalnih resursa, učinkovita naplata njihovog korištenja, te dinamička i bez prekida u radu aplikacije preraspodjela računalnih resursa. Bez obzira na očite prednosti Cloud Computing okruženja za HPC, trenutni omjer korištenja HPC aplikacija u računalnim oblacima u odnosu na one koje se izvode u tradicionalnom HPC okruženju još je uvijek vrlo malen. Neki od razloga za takvo stanje su očiti, dok kod drugih nije tako. Tako na primjer tradicionalni proizvođači HPC opreme i aplikacija još uvijek pokušavaju iskoristiti njihova ulaganja u postojeća rješenja na način koji pogoduje tradicionalnom načinu rada HPC-a. Nadalje, iako je prisutan snažan razvoj virtualizacijske tehnologije, još uvijek postoje otvorena pitanja, poput skaliranja HPC aplikacija u virtualnom okruženju, te suradnje između fizičkih i virtualnih resursa. Tu se također nameće i pitanje podrške standardnih HPC alata za uspostavu virtualiziranih okolina. Na kraju, ali ne i manje važno, pojava heterogenog HPC računalstva koje se zasniva na kombinaciji standardnih procesora (CPU) i onih specijalizirane namjene (GPU procesori i akceleratori) nameće izazovna pitanja, poput skalabilnosti, učinkovitog Linpack mjerenja performansi, te razvoja HPC aplikacija u heterogenom CPU/GPU okolišu. Uz sve to, unatoč tome što mnogi proizvođači HPC opreme i aplikacija govore o tome kako imaju potpuno funkcionalna HPC rješenja u računalnim oblacima, postoji potreba da se osigura odgovore i na dodatna pitanja: Jesu li potpuno ispunjena obećanja mogućnosti virtualizacije svih komponenti sustava, ne samo onih virtualnih, nego i fizičkih? Koje vrste računalnih oblaka su podržane: privatne, javne ili njihova kombinacija, tj. hibridni oblaci? Kako dobro pojedina HPC aplikacija skalira na nekom računalnom oblaku? U pokušaju odgovora na postavljena pitanja, u ovom radu je dat pregled trenutnih HPC rješenja u oblacima. Osim toga, namjera ovog rada je da posluži kao vodič za onoga tko je spreman napraviti pomak od standardnih HPC rješenja prema onima u Cloud Computing okruženju. Naposljetku, da bi se što je moguće više olakšao izbor HPC Cloud rješenja, predložena je njihova klasifikacija u tri kategorije, od kojih svaka reflektira područja i načine primjene pojedinih rješenja, stupanj podržanosti virtualizacije, te, zbog uglavnom javno nedostupnih rezultata mjerenja performansi, do sada iskazanu HPC ekspertizu proizvođača tih rješenja.

**Ključne riječi:** *pregled HPC Cloud rješenja; prvi, drugi i treći rang HPC Cloud rješenja; budući razvoj i trendovi u HPC Cloud rješenjima*

## 1 Introduction

There is an increasing interest of doing HPC in cloud computing environments. Cloud computing promises to effectively deliver resources on demand and when needed, drive IT costs down, maximize utilization of IT infrastructure and applications [1]. An ultimate goal is to bring HPC functionality to the desktop of any researcher, despite a fact that he is sitting in a HPC datacenter, or exploiting its resources while connected to them via Internet connection. In case local resources are exhausted, it should be possible to grab additional HPC resources from external public Cloud, quickly and without interrupting computing jobs. Further on, the portal from which a typical HPC user (at most the times not very skilled IT person) can deploy his HPC application should be simple, intuitive and easy to use. The benefits of doing HPC in Cloud environments are obvious [2], and a list of hardware and software vendors offering HPC Cloud solutions is steadily growing, as for example shown in [3]. However, and because real case references are still

mostly missing, these solutions need to be carefully examined case by case. On the other side, there is still a need for HPC cloud enabled applications. For example, with an advent of HPC Grids [4], it quickly becomes clear that applications need to be aware [5] about the Grid topology and its performance properties, otherwise Grid middleware [6] will not be able to balance application workload in a most effective way. Additionally, in opposite to the standard HPC benchmarks like LINPACK, STREAMS and NAS Grid benchmark [7], there is no standardized and widely adopted HPC Cloud benchmark available as of today. In other words, we have no widely adopted metric to measure the effectiveness of a certain HPC Cloud system, except for using existing HPC benchmarks [8]. While existing benchmarks may work fine for performance evaluation of a certain HPC Cloud, they will fail to measure some exclusive Cloud properties, for example, the speed of cloud bursting process and provisioning of infrastructure and applications. Most Cloud system providers are allowing an access to virtual computing resources, by exploiting

technologies based on hypervisor paradigms, like Microsoft's Hyper-V, vSphere from VMWARE, or Xen from open source community. In opposite, very few of them are able to provide hardware virtualization, like Cisco is able to do with its virtual switch technology on the network layer, or Hewlett-Packard, providing virtualization on all infrastructure levels, i.e. server, storage and network. And handling equally both virtualized and physical resources is of the utmost importance for doing HPC computations effectively, because many HPC applications are not only hungry for processing resources like CPUs and GPUs, but generate a large amount of data as well [9]. The movement of data in HPC Cloud occurs not only within a datacenter, but between geographically dispersed datacenters, too. Therefore, the most effective way is to leave the data where it is, and logically move physical servers instead. Because the logical move of the physical server implies the movement of its properties, i.e. its MAC addresses, WWN addresses and his working profile, to some other physical server, it is necessary to implement switches with a virtualization layer, which will preserve these addresses on a new server. There is another big challenge in HPC Cloud, and this is load balancing. An average HPC environment has to be precisely load balanced, sometimes very rarely to allow only one HPC application to privilege all available resources and to run on them as fast as possible, and most of the time supporting the execution of a bunch of smaller HPC batch jobs running in a certain HPC cluster. The widely adopted strategy in the industry standard HPC job schedulers like Moab from Adaptive Computing [10] and LSF from Platform Computing [11] is to utilize some of nodes as much as possible, by still allowing jobs to run smoothly and effectively. Other nodes in a cluster should be ideally free from jobs, thus allowing system administrator or automatic procedures to put them in a low-power consumption mode, and when needed, to shut them down or restart. While it is relatively simple to balance a load in a HPC cluster with all servers of equal performance and interconnects with equal bandwidth, it is not the same case when doing the same in the heterogeneous Cloud environment [12]. Even defining a Cloud is not a simple task [13]. In addition, some vendors with no HPC expertise, like VMWARE, are providing key components for HPC Cloud solutions, while other, traditional HPC companies without significant Cloud expertise, like Platform and Adaptive Computing, are providing components of the essential importance for any HPC Cloud solution. Besides, there are vendors delivering all three standard models of HPC Cloud, i.e. public, private and hybrid, while others, like Google and Amazon, provide only public HPC Clouds. Even in this area, there are differences; while formerly mentioned companies provide HPC as a service on their public Clouds that are general computing platforms, others, like Penguin, provide HPC as a service on its specialized public HPC Cloud platforms. Therefore, we will take the following approach to review HPC Cloud solutions currently available on the market: we will split HPC Cloud solutions in three tiers. We will define these tiers as follows:

First tier HPC Cloud solutions are complete solutions, supporting private, public and hybrid HPC Clouds. Additionally, they provide the support both for physical and virtual resources.

Second tier HPC cloud solutions are complete solutions as well, but either not supporting all classes of Clouds or not supporting both physical and virtual resources.

Third tier HPC cloud solutions are components only, but of the key importance for the first and the second tier HPC Cloud solutions.

The proposed tiring of HPC Cloud does not mean that the solution belonging for example to the first tier is better than the solution in the second tier. We classified HPC Cloud solutions according to this scheme because of a better overview of the current HPC Cloud solutions. Besides, it supports standard model of Cloud computing. For example, a high-quality job scheduler, as a third tier HPC Cloud solution, can deliver much more benefit to HPC user than a complete tier one solution, providing a self-service portal, charge back capabilities, and support for both virtual and physical environments, but expressing performance and scalability issues. Performance is at the heart of HPC, and that holds true for HPC Cloud environment as well.

## 2 An overview of HPC Cloud solutions

At this moment, we are not aware of any first tier HPC cloud working solution. However, by looking at the publicly available data, we see several companies plus one coalition offering all three classes of Clouds, i.e. private, public and hybrid, and providing support for both physical and virtual resources. These companies are, in the alphabetic order: Hewlett-Packard, IBM and Microsoft, plus VCE, a coalition of Cisco, EMC and VMWARE. Their Cloud solutions are platforms designed to run every application including HPC. However, no publicly available benchmarking results are available, which could prove these solutions are able to deliver performance and scalability needed for HPC. What we have are HPC references of these vendors, and among these companies, VCE is still missing them. That does not mean VCE Cloud solution is not good enough for HPC. It expresses our following view: if someone has not HPC references or related benchmarking data, we are not able to validate his solution.

### 2.1 First tier HPC Cloud solutions

Hewlett-Packard delivered its first Cloud solution in 2009, and named it HP BladeSystem Matrix [14]. At its introduction, it was a private Cloud solution only, able to provision physical and virtual servers, storage systems and network. In the meantime, HP BladeSystem Matrix has evolved and has been renamed to HP CloudSystem Matrix. A year later, Hewlett-Packard announced one-step further towards the completeness of their Cloud system offering, and called it CSA (Cloud Server Automation) for BladeSystem Matrix. CSA for BladeSystem Matrix is a software suite that, besides infrastructure, enables provisioning of applications. As a next step, Hewlett-Packard renamed BladeSystem Matrix to CloudSystem

Matrix, and Cloud Server Automation to Cloud System Automation. Relying on HP Blade server and Virtual Connect technology, CSA for BladeSystem Matrix combined with HP CloudSystem Matrix delivers Cloud solution that provides both Infrastructure-as-a-Service (IaaS) and Software-as-a-Service (SaaS). Hewlett-Packard is calling this combination CloudSystem Base. For hybrid and public Cloud support, company has two additional Cloud offerings, both with HP CloudSystem Matrix and HP CSA for CloudSystem Matrix as a base: HP CloudSystem Standard, which, on the top of HP CloudSystem Matrix with CSA, enables additional functionality of bursting resources between private and public Clouds, plus additional, web based Self Service Portal. For someone that needs a public Cloud solution, HP delivers its CloudSystem Enterprise solution, with a full functionality of a HP CloudSystem Standard, plus additional support for third party hardware and software, and network automation at a pure physical layer, needed for public Clouds supporting a very large number of users.

HP CloudMaps are additional components allowing automatic and standardized provisioning of applications on HP CloudSystem. HP CloudMaps are tested software templates, which provide immediate provisioning of Oracle, SAP, McAfee and other industry standard solutions. Currently, we do not see HPC related CloudMaps available. Therefore, users willing to deploy a certain HPC application on a HP CloudSystem, needs to build its own CloudMaps. Last but not least, HP has offered this year its own public solution, built on Linux Ubuntu engine. This engine is well known among Open Source community as a scalable and highly reliable, thus offering someone to put its HPC applications on it with no constraints.

IBM is a company that has as well its proprietary blade technology, and with an advent of Cloud computing, delivered its Cloud system, called IBM CloudBurst [15]. The goal was the same as in the case of HP BladeSystem Matrix, i.e. to fulfill automatic provisioning of servers, network and storage infrastructure. Because the provisioning of hardware only is not a fully functional cloud solution, IBM recently announced integration of IBM CloudBurst system with WebSphere CloudBurst Appliance, thus enabling automatic provisioning of applications as well. Enhancing IBM CloudBurst with the IBM WebSphere CloudBurst Appliance provides a fully functional Cloud platform. With WebSphere CloudBurst Appliance, one can develop and deploy HPC applications in CloudBurst environment. Similar to what Hewlett-Packard is doing with its CSA for BladeSystem Matrix, IBM CloudBurst solution with WebSphere CloudBurst Appliance supports equally physical and virtual nodes. On the other side, IBM CloudBurst solution has ability to integrate with the equipment outside the single datacenter, thus providing a support for public, private and hybrid Clouds as well. Very recently, IBM announced next generation cloud services, aka IBM SmartCloud services, which increase performance and availability. These services are built on IBM's existing public cloud. Primarily designed for business applications, there seem to be no obstacles to extend these services as a platform for doing HPC in the

cloud. Both Hewlett-Packard and IBM have a strong expertise in HPC and share most of the entries in <http://www.top500.org/>, the list of most powerful Supercomputers today. Therefore, we expect that they are able to implement high-quality HPC solutions in their Cloud offerings.

On the contrary, Microsoft started to play in a HPC market as a niche player for relatively small HPC installations, only several years ago. However, the deployment of HPC solutions based on their Windows HPC Server 2008 [16] is rumpling at a fast rate, and as of today, Microsoft has received a lot of attention as a HPC solution provider. Former limitation on the number of computing nodes vanished, and now it is possible to build a Windows based HPC cluster on hundreds of nodes and thousands of cores. Windows HPC Server 2008 R2 is a base for running HPC applications and integrated solution offering all HPC middleware key components needed to run a certain HPC cluster efficiently: application launcher, MPI, job scheduler, cluster management and reporter. What is also important is its ability to integrate with Microsoft Azure platform [17], and to burst on demand and when needed additional resources from it. On the other side, there are some limitations still present: a scaling-out in HPC clusters with a very large number of computing nodes, and a missing support for heterogeneous clusters consisting of Windows and Linux computing nodes.

Although it is possible to build and manage HPC clusters of small and medium sizes without job schedulers, it is hard to obtain an effective job scheduling in a large HPC cluster. Microsoft is the only company among the first tier HPC Cloud vendors providing complete HPC solutions without a third party job scheduler. Proprietary job scheduler is an integral part of their Windows HPC Server 2008 solution. Contrary to the exclusive and small first tier club, there are many more vendors and their HPC Cloud solutions in the second tier. We will not give a review of all HPC Cloud solutions in this tier available. Our focus will be on the vendors and solutions with the greatest presence on HPC Cloud market. We will review HPC Cloud solutions from Amazon, Google, Penguin, R Systems, SGI and Univa.

## 2.2 Second tier HPC Cloud solutions

Amazon is present on HPC Cloud Computing market with its EC2/S3 infrastructure [18]. It is a public Cloud, providing server and storage resources on-demand, and Amazon is charging the usage of its resources at acceptable rates. It is maybe the main reason why EC2 based HPC Cloud solutions have undergone a good benchmarking process. Research performed in [19] showed that the performance of EC2/S3 Cloud is bellow HPC dedicated clusters. It is due to the inter-communication of nodes with more processing power and those with a less performance power in large clusters. However, for a small to medium HPC clusters, Amazon EC2/S3 Cloud platform provides a reasonable performance/price solution.

Google infrastructure together with its WebApps [20] provides a public HPC Cloud in a similar way Amazon is doing. Its Cloud engine uses a Map-Reduce algorithm

[21], and performs well within nodes of equal performance. When deploying a certain HPC application on heterogeneous nodes, performance degradation is likely to occur.

Penguin Computing is a company providing HPC private Cloud solutions [22]. It offers integrated hardware and software HPC Cloud solution, and the key component is Scyld ClusterWare, a Linux based clustering software that makes a cluster appear and act like a single system. With added self-service portal, HPC user has a platform able to execute HPC applications in the Cloud look-and-feel manner, preserving effective control and management of large number of nodes. The final confirmation of their already acquired status of HPC experts is a very recent contract with the University of Virginia, which allows them to sell hardware, software and services to member institutions.

R Systems takes a different approach. It sells software that brings Cloud features, like self service portal and effective charge backing, in traditional HPC data centers [23]. In that way, it is a solution delivering standard HPC Cloud functionality in traditional HPC environments. The company partnered this year with Wolfram research, enabling its Mathematica to run on R Smart Cluster as a HPC service. Nevertheless, R Systems already proved their ability to deliver valuable HPC solution, by building the 44th fastest supercomputer on the TOP500 list.

SGI, a company with a long HPC expertise and references, decided to put a foot in a demanding HPC Cloud market. It sells a solution under the name SGI Cyclone [24]. This solution is a public HPC cloud, and it runs on SGI standard HPC hardware platforms: SGI Shared Memory Compute Cluster with up to 64 cores, and SGI Distributed Memory Compute Cluster with up to 256 cores. On top of it, there is a self-service portal from Nimbis Services Inc., available on Nimbis website. From this portal, user can request either HPC infrastructural resources, or HPC applications. There are eleven HPC applications on SGI Cyclone available, among others OpenFOAM, NAMD and mpiBLAST. Currently SGI Cyclone is available in US only.

Univa delivers UniCloud solution [25], a management software product that delivers capabilities for enabling traditional infrastructure to private, public and hybrid HPC Cloud computing environments. It supports a build of private virtual Clouds based on Oracle VMs and VMWARE, and supports several public Clouds, i.e. Amazon AWM, Rackspace and GoGrid. UniCloud delivers bursting capabilities between private and public Clouds, thus supporting hybrid HPC Clouds as well. The latest release of Univa HPC Cloud solution software, UniCloud 3.0, delivers further enhancements in scalability, manageability and availability areas of HPC solutions running in public clouds.

### 2.3 Third tier HPC Cloud solutions

Companies traditionally being involved in HPC business started to deliver Cloud aware HPC applications. For example, Mathworks has standardized usage of its parallel MATLAB library on Amazon's EC2 Cloud [26]. While it is likely that HPC application vendors will

deliver more attention to the application scalability in heterogeneous environments, there is still a lack of the effort in this area. Another issue with HPC applications in the Cloud is that they need to have APIs to Cloud self-service portal, if someone likes to provision them in a standardized way.

In addition, we could expect more development in the HPC development tools. For example, OpenMPI [27], a standard language for HPC programming, is still missing support for virtual machines.

On the other side, a good attitude comes from companies delivering traditional and key HPC components. While these companies still develop their HPC job schedulers, they start to invest more in Cloud technologies. Adaptive Computing with its Moab Adaptive Computing Suite [28] delivers automatic resources provisioning, optimization, chargeback, management and reporting in HPC Cloud environments.

Platform Computing is another traditional HPC oriented company that recently started to support HPC Cloud environments. This company provides two approaches to support clouds: Private Cloud solutions for IT enterprise data centers and Cloud bursting for HPC [29]. Platform ISF is a product that provides management of the application workloads across multiple VM technologies and provisioning tools. Platform LSF and Platform MultiCluster are products that enable Cloud bursting from external resources.

The third company we will mention here is VMWARE. Although traditionally not a HPC company, its virtual machines are now the components of many HPC Cloud solutions. Besides Xen and Microsoft's Hyper-V, its vSphere product [30] provides a high quality virtualization layer in HPC Cloud environments [30].

### 3 Future trends and needs

Just now HPC cloud solutions are rare compared to traditional HPC systems, and today they are mainly present as HPC public clouds. For example, based on IDC report [31], an expectation is that by the year 2015 HPC public clouds will represent less than 5 % of all server CPU hours within HPC. It seems that for the real boost of HPC Cloud, the wide collaboration between Cloud providers, HPC application vendors and HPC community is critical. We expect that both HPC public and private clouds will grow in the next years. HPC private cloud solutions will be dominant for organizations needing a more collaborative work on the same datasets, and HPC public solutions as a world-wide working environment for a smaller research entities, institutions and companies not willing or being able to invest in HPC clusters that will be deployed rarely. In addition, we expect that standardized HPC Cloud benchmarks will help to solve the main question mark, i.e. scalability of HPC cloud solutions. While classical HPC benchmarks are still appropriate for HPC private cloud solutions, for the HPC public and hybrid solutions they are not. The main reason is their orientation towards performance measurement of unified and homogeneous HPC clusters. Good candidates for this are HPC Grid benchmarks, because they allow performance measurement of geographically dispersed and heterogeneous HPC clusters.

Moreover, for a wider acceptance of a HPC cloud solution, it is crucial that main cloud providers, like Amazon, Google, Hewlett-Packard, IBM, Microsoft and VMWARE start to invest more in research, development and integration between their cloud operating environments and the standard HPC software coming either from the commercial vendors or from open source community. Our view is that, despite the publicly available LINPACK performance data, real petaflops Supercomputer is still a vision and hard to obtain, pushing technological and application design limits behind these available today. We expect that the generation behind, able to execute hexaflops, will be on Cloud.

#### 4 Conclusion

Doing HPC in a cloud is a still challenging task, deserving the highest attention, both from HPC and Cloud research community. While the common agreement is that Cloud will boost HPC to the next level regarding performance and availability, there are still questions that need answers. One of the most crucial among them is a scalability of HPC on the Cloud and the performance measurement of HPC applications executed on Cloud. Lessons learned from Grid Computing show that any HPC application running on HPC cluster within one datacenter is likely not to scale linearly when deploying it on several HPC clusters in geographically dispersed datacenters. The reasons are several, from those regarding the overhead of the management software, up to the optimal workload scheduling of the jobs on the nodes/datacenters with various performance functionality, and interconnected with links of different bandwidth. It is also of the utmost importance to classify and measure the performance of certain HPC cloud systems. As our review shows, public HPC Clouds dominate today's HPC Cloud market, like Amazon's EC2/S3 and Google's WebApps. HPC private clouds still need to compete against standard HPC clusters, and to prove their benefits regarding ease of use, availability and scalability. Our review of HPC cloud solutions also shows that very few companies are able to provide a full PPH (Private Public and Hybrid) model of HPC Cloud. The reason is that the HPC Cloud solution still requires a lot of expertise, both in HPC and Cloud areas.

On the other side, and because the benchmarking data are still mostly missing, we were unable to compare the reviewed HPC solutions in terms of their performance and scalability. User willing to deploy HPC Cloud solution has several options: either to run it entirely in a private or public Cloud, or to combine these two approaches in a hybrid HPC Cloud model. Because of the still lacking standards, a careful evaluation of the Cloud bursting capabilities between private and public Clouds should take place. Otherwise, it may happen that the bursting between private and public HPC Clouds from various vendors does not perform well. Another interesting option for running HPC in a Cloud is the transformation from the traditional HPC to HPC Cloud. Our review has shown that there is a bunch of companies providing this possibility. This option is of the utmost importance for the organizations already in HPC, because in that way they

can protect investment and gain Cloud functionality in their HPC environment.

#### 5 References

- [1] Vecchiola, C.; Pandey, S.; Buyya, R. High-Performance Cloud Computing: A View of Scientific Applications, Proceedings of the 10th International Symposium on Pervasive Systems, Algorithms and Networks (I-SPAN 2009, IEEE CS Press, USA), Kaohsiung, Taiwan, December 14-16, 2009.
- [2] Nitzberg, B. Ph. D. CTO, PBS GridWorks Altair Engineering, Inc. Aug 2009, High Performance Computing with Clouds: Past, Present, and Future, <http://www.linuxworldexpo.com/storage/10/documents/CI2%20Bill%20Nitzberg.pdf>.
- [3] Cloud Computing vendors for HPC, [www.hpcwire.com/hpcwire/2009-11-03/cloud\\_computing\\_vendors\\_for\\_hpc.html](http://www.hpcwire.com/hpcwire/2009-11-03/cloud_computing_vendors_for_hpc.html), November 03, 2009.
- [4] Stockinger, H. Defining the Grid: A Snapshot on the Current View, *The Journal of Supercomputing*, 42, 1(2007), pp. 3-17.
- [5] Herrera, J.; Huedo, E.; Montero, R. S.; Liorente, I. M. Developing Grid-Aware Applications with DRMAA on Globus-based Grids, Technical Report, 2008.
- [6] Foster, I. Globus Toolkit Version 4, Software for Service-Oriented Systems, IFIP International Conference on Network and Parallel Computing, Springer-Verlag LNCS 3779, pp. 2-13, 2005.
- [7] Snaveley, A.; Chun, G.; Casanova, H.; Wijngaart, R.; Frumkin, M. Benchmarks for grid computing: a review of ongoing efforts and future directions, // *ACM SIGMETRICS Performance Evaluation Review*, 30, 4(2003), pp. 27-32.
- [8] Evangelinos, C.; Hill, C. Cloud Computing for Parallel Scientific HPC Applications: Feasibility of Running Coupled Atmosphere Ocean Climate Models on Amazon EC2, Proceedings of CCA-08, ACM, 2008.
- [9] Mergen, M.; Uhlig, V.; Krieger, O.; Xenidis, J. Virtualization for high-performance computing. // *ACM SIGOPS Operating Systems Review*, 40, 2(2006), pp. 8-11.
- [10] <http://www.adaptivecomputing.com/products/moab-hpc.php>
- [11] <http://www.platform.com/workload-management/high-performance-computing>
- [12] Xie, J.; Yin, S.; Ruan, X.; Ding, Z.; Tian, Y.; Majors, J.; Manzanares, A.; Qin, X. Improving MapReduce performance through data placement in heterogeneous Hadoop clusters. // Proceedings of the 24th IEEE International Parallel & Distributed Processing, Workshops and Phd Forum (IPDPSW), pages 1 – 9, Atlanta, GA, 19-23 April 2010.
- [13] Armbrust, M.; Fox, A.; Griffith, R.; Joseph, A. D.; Katz, R. H.; Konwinski, A.; Lee, G.; Patterson, D. A.; Rabkin, A.; Stoica, I.; Zaharia, M. Above the Clouds: A Berkeley View of Cloud, Computing. Technical Report No. UCB/EECS-2009-28, February 10, 2009.
- [14] <http://h18004.www1.hp.com/products/blades/components/matrix/main.html>
- [15] <http://www-01.ibm.com/software/tivoli/products/cloudburst/>
- [16] <http://www.microsoft.com/hpc/en/us/default.aspx>
- [17] [http://www.microsoft.com/windowsazure/Whitepapers/HP\\_CServerAndAzure/default.aspx](http://www.microsoft.com/windowsazure/Whitepapers/HP_CServerAndAzure/default.aspx)
- [18] <http://aws.amazon.com/ec2/>
- [19] Edward Walker, Benchmarking Amazon EC2 for high-performance scientific computing, [www.usenix.org/publications/login/2008-10/openpdfs/walker.pdf](http://www.usenix.org/publications/login/2008-10/openpdfs/walker.pdf)
- [20] <http://appengine.google.com>

- [21] Jeffrey, D.; Sanjay, G. MapReduce: Simplified Data Processing on Large Clusters Dean, In Proceedings of OSDI 2004, San Francisco, CA, 2004.
- [22] <http://www.penguincomputing.com/POD>
- [23] <http://www.rsystemsinc.com/>
- [24] [http://www.sgi.com/products/hpc\\_cloud/cyclone/](http://www.sgi.com/products/hpc_cloud/cyclone/)
- [25] <http://www.univa.com/products/unicloud-tech-specs.php>
- [26] Parallel Computing with MATLAB on Amazon Elastic Compute Cloud, [http://www.mathworks.com/programs/techkits/ec2\\_paper.html](http://www.mathworks.com/programs/techkits/ec2_paper.html), technical whitepaper
- [27] <http://www.open-mpi.org/>
- [28] <http://www.adaptivecomputing.com/products/moab-adaptive-computing-suite.php>
- [29] <http://www.platform.com/private-cloud-computing/clouds#privatecloud>
- [30] <http://www.vmware.com/products/vsphere/overview.html>
- [31] <http://www.marketresearch.com/product/display.asp?productid=2640009&xs=r>

#### **Authors' addresses**

##### ***Draško Tomić***

Hewlett-Packard Croatia  
Radnička cesta 41  
10000 Zagreb, Croatia  
E-mail: [drasko.tomic@hp.com](mailto:drasko.tomic@hp.com)

##### ***Dario Ogrizović***

University of Rijeka  
Faculty of Maritime Studies  
Center for advanced computing and modeling  
Studentska ulica 2  
51000 Rijeka, Croatia  
E-mail: [dario@cnrm.uniri.hr](mailto:dario@cnrm.uniri.hr)

##### ***Zlatan Car***

Faculty of Engineering  
University of Rijeka  
Vukovarska 58  
51000 Rijeka, Croatia  
E-mail: [car@riteh.hr](mailto:car@riteh.hr)