

Cloud Computing Research

http://www.pds.ewi.tudelft.nl/~iosup/research_cloud.html

Rationale

why and how is this work relevant?

Cloud computing is an emerging commercial infrastructure paradigm that promises to eliminate the need for maintaining expensive computing hardware. Through the use of virtualization and resource time-sharing, clouds address with a single set of physical resources a large user base with different needs. Thus, clouds promise to enable for their owners the benefits of an economy of scale and, at the same time, reduce the operating costs for many applications. For example, clouds may become for scientists an alternative to clusters, grids, and parallel production environments.

People

who is part of the group?

- **Undergraduate Students:** Adrian Lascateu (Politehnica University of Bucharest, Romania).
- **Graduate Students:** **Nezih Yigitbasi** (TU Delft), Vlad Nae (U. Innsbruck, Austria).
- **Collaborators:** **Dick Epema**, **Henk Sips** (both TU Delft), Thomas Fahringer, Radu Prodan (both U. Innsbruck, Austria), Nicolae Tapus (Politehnica University of Bucharest, Romania).
- **Yours truly:** **Alexandru Iosup**.

Main Research Questions

what do we try to achieve?

1. **What is the actual performance of virtualized cloud resources?** Many production clouds, including the largest publicly-accessible commercial clouds such as the Amazon Web Services and the Google App Engine, use virtualized resources to address diverse user requirements with the same set of physical resources. Virtualization can introduce performance penalties, either due of the additional middleware layer or to the interaction of workloads belonging to different virtual machines. Do virtualized resources deliver the same performance regardless of the application? In particular, are scientific applications affected by execution on virtualized resources?
2. **What guarantees do we have about the good performability of clouds over long periods of time?** A major impediment to cloud adoption at large is their perceived instability, due, in lack of hard evidence, to novelty ("clouds are a technology too immature to be reliable"). Even if a cloud is available and works well today, it may well happen that it will not tomorrow. Does performance change over time (for the worse)? Are clouds really available all the time?
3. **Which new applications can make use of clouds?** Commercial clouds are new to the public. What applications that we could not previously afford to run are now enabled by clouds? What applications can function well under the availability and performance profiles of the current production cloud services?

Main Achievements

what did we do?

1. **Evaluated the performance of resources from four production, commercial clouds.** We have added to GrenchMark the C-Meter tool for evaluating the performance of cloud resources [3]. We have studied [2,6,9] the performance of resources from four production, commercial clouds: Amazon Elastic Compute Cloud (EC2), Mosso, Elastic Hosts, and GoGrid.
2. **Evaluated the variability of the performance delivered by production cloud services.** We have collected year-long traces and, based on them, studied [8] the performance of over fifteen operations provided by nine services in two clouds, Amazon Web Services and Google App Engine.
3. **Designed CAMEO, a framework for continuous analytics for massively multiplayer online games (MMOGs) using cloud resources.** We have provided [4] a first estimate of the feasibility and costs of performing continuous analytics for MMOGs on cloud resources.
4. **Within edutain@grid, analyzed the feasibility of running massively multiplayer online games (MMOGs) on cloud resources.** We have analyzed [1,7] the potential gains of running MMOGs on cloud resources, based on an ideal cloud/data center model. We have started to investigate [5] the impact of virtualization on running MMOGs on cloud resources.

Main Findings


what did we find?

1. "The performance of the resources leased by four production clouds, including Amazon EC2, for running scientific applications is at the moment below the theoretical peak and below the maximum efficiency obtained in other specialized environments." [2,6,9]
2. "The costs of running grid workloads in clouds may be much lower than the costs of building and maintaining the grids that can run these workloads." [9]
3. "The resource allocation strategy employed by the cloud user can be insignificant or can lead to 10-30% cost increase, when running long-term, grid-like workloads on clouds." [2,9]
4. "The performance of the production cloud services offered by Amazon Web Services and Google App Engine varies over time, and can alter significantly the performance and cost profiles of large-scale applications such as job execution in scientific computing, virtual goods trading in social networks, and state management in social gaming." [8]
5. "Continuous analytics for MMOGs on cloud resources is feasible and shows good promise of being cost-effective." [4]
6. "Dynamic resource allocation from clouds can lead to a tenfold reduction of the platform operation costs for massively multiplayer online games (MMOGs)." [1,7]

Publications

[journals/conferences/workshops](#) | [all PDS group publications](#) | [my publications \(with BibTeX\)](#) | [my DBLP entry](#) | [my ACM DL entry](#)

2010 [details](#) <-- [click to see more details](#)

- [9] **A. Iosup**, S. Ostermann, N. Yigitbasi, R. Prodan, Th. Fahringer, and D. Epema, *Performance Analysis of Cloud Computing Services for MTC-Based Scientific Computing*, (under submission). 


keywords cloud computing, scientific computing, performance evaluation.

info Evaluated the performance of four cloud computing providers (Amazon EC2, Mosso, Elastic Hosts, and GoGrid) when executing scientific computing micro-benchmarks and applications (HPCC, Imbench, Bonnie, CacheBench). Investigated through trace-based simulations the use of clouds vs. other scientific computing infrastructures such as grids and parallel production environments.

- [8] **A. Iosup**, N. Yigitbasi, and D. Epema, *On the Performance Variability of Production Cloud Services*, (under submission).


keywords cloud computing, performance variability, social applications, Amazon Web Services, Google App Engine.

info Evaluated the performance variability of over ten cloud service operations from two major production clouds (Amazon Web Services and Google App Engine). Investigated through trace-based simulations the impact of the variability observed for the studied cloud services on three large-scale applications, job execution in scientific computing, virtual goods trading in social networks, and state management in social gaming.

- [7] R. Prodan, V. Nae, and **A. Iosup**, *Dynamic Resource Provisioning in Massively Multiplayer Online Games*, In IEEE Transactions on Parallel and Distributed Systems (TPDS), 2010 (accepted, in print). 

keywords massively multiplayer online games, platform, virtualization, cloud computing.

info (Improved version of our SC|08 article) We have investigated a new dynamic resource provisioning method for MMOG operation using external data centres as low-cost resource providers. We have introduced a comprehensive MMOG load model that takes into account both the player interaction type and the population size. We have assessed using trace-based simulation the impact of the data centre policies on the quality of resource provisioning. Last, we have presented experimental results showing the real-time parallelization and load balancing of a real game prototype using data centre resources provisioned using our method and show its advantage against a typical client threshold approach.


- [6] S. Ostermann, **A. Iosup**, N. Yigitbasi, R. Prodan, T. Fahringer, and D. Epema, *A Performance Analysis of EC2 Cloud Computing Services for Scientific Computing*, In D.R. Avresky et al. (Eds.): Cloudcomp 2009, LNICST 34, pp. 115–131, 2010. 

keywords cloud computing, scientific computing, performance evaluation.

info (Cleaned version of TU Delft Technical Report PDS-2008-006) Evaluated the performance of one cloud computing provider (Amazon EC2) when executing scientific computing micro-benchmarks and applications (HPCC, Imbench, Bonnie, CacheBench). Identified improvements for scientific computing on cloud resources.

[click to close](#) [details](#)

2009 [details](#) <-- [click to see more details](#)

- [5] V. Nae, **A. Iosup**, R. Prodan, and T. Fahringer, *The Impact of Virtualization on the Performance of Massively Multiplayer Online Games*, In *ACM/IEEE Symposium on Network and Systems Support for Games (NetGames 2009)*, Paris, France, November 23-24, 2009 (accepted, acceptance ratio **27%**=10/37). 

keywords massively multiplayer online games, platform, virtualization, cloud computing.

info We have proposed a new hybrid resource provisioning model that uses a smaller and less expensive set of self-owned data centers, complemented by virtualized cloud computing resources during peak hours. Using real traces from RuneScape, one of the most successful contemporary MMOGs, we have evaluated with simulations the effectiveness of the on-demand cloud resource provisioning strategy for MMOGs.

- [4] A. Iosup, *CAMEO: Continuous Analytics for Massively Multiplayer Online Games on Cloud Resources*, In 2nd International Workshop on Real Time Online Interactive Applications on the Grid (ROIA 2009), Delft, the Netherlands, August 24, 2009, Springer, LNCS vol.?, p.1--10 (accepted, in print).



keywords massively multiplayer online games, analytics, cloud computing.

info We have introduced CAMEO, an architecture for Continuous Analytics for Massively multiplayer Online games on cloud resources.

- [3] N. Yigitbasi, A. Iosup, S. Ostermann, and D.H.J. Epema, *C-Meter: A Framework for Performance Analysis of Computing Clouds*, In the International Workshop on Cloud Computing (Cloud 2009), May 18-21, 2009, in conjunction with CCGrid'09 (accepted).



keywords cloud computing, GrenchMark, performance evaluation, C-Meter.

info We extend GrenchMark, our framework for performance evaluation in large-scale distributed systems, with C-Meter, a component to assess the performance of computing clouds.

click to close [details](#)

2008 [details](#) <-- click to see more details

- [2] S. Ostermann, A. Iosup, N. Yigitbasi, R. Prodan, T. Fahringer, and D. Epema, *A Performance Analysis of EC2 Cloud Computing Services for Scientific Computing*, TU Delft Technical Report PDS-2008-006, Dec 2008.



keywords cloud computing, performance evaluation, performance analysis, benchmarking, Bonnie, LMBench, CacheBench.

info Evaluated the performance of one cloud computing provider (Amazon EC2) when executing scientific computing micro-benchmarks and applications (HPCC, lmbench, Bonnie, CacheBench). Identified improvements for scientific computing on cloud resources.

- [1] V. Nae, A. Iosup, S. Podlipnig, R. Prodan, D.H.J. Epema, T. Fahringer, *Efficient Management of Data Center Resources for Massively Multiplayer Online Games*, In the ACM/IEEE SuperComputing Conference on High Performance Networking and Computing (SC'08), Nov 10-16, 2008 (accepted; acceptance ratio **20%**).



keywords massively multiplayer online games (MMOG), cloud computing, data centers, virtual environments, resource provisioning, games.

info We have analyzed through trace-based simulations the potential of running MMOGs entirely in clouds and data centers. We modeled ideal clouds and data centers, and found a theoretically achievable ten-fold reduction in the resource consumption.

click to close [details](#)

File types: .bib | .pdf | .ps | .doc | .ppt | .tgz | .zip | description | any
Other information: Electronic Edition | Conference/Journal ToC (DBLP/ACM DL)

Last modified: Wed, 27 January, 2010 12:20 PM
The newest version of this page can be found at: http://www.pds.ewi.tudelft.nl/~iosup/research_cloud.html