

## Overview

Cloud computing is probably the hottest topic in current Internet systems. The cloud brings the illusion of a virtually infinite computing infrastructure/platform, where the immediate scaling of resources, the employment of advanced billing mechanisms allowing for a pay-per-use model on shared multitenant resources, the simplified programming mechanisms (platform), etc. are among the most relevant features.

This is a new model for managing applications so that they can scale when it is needed. Indeed, cloud-promised scalability is considered what makes any particular service something more than “just an outsourced service with a prettier marketing face”. As a result of its relevance, the wealth of systems dealing with “cloud application scalability” is slowly gaining weight in the available literature. Also industrial systems such as Amazon Autoscale, RightScale, Google App Engine, etc. deal with the scalability of cloud applications, platforms and data.

While the literature in infrastructure or data scaling is extensive, a lot is still to be done when it comes to scaling both the virtual resources built from the barebones infrastructure of the datacenter (e.g. I/O access of VMs, scalable networks controlled by the user, etc.) and the services provided to application developers (such as databases, container systems, etc.). Of course, this scaling potential needs to be available without imposing a huge burden for application development itself. Indeed, the goal would be to release developers from having to control scalability-related topics (e.g. thread control, data replicas management, etc.) due to the usage of cloud-dwelling services.

If new cloud architectures, along with their companion tools and middleware, are able to meet those goals, cloud technologies have the potential to change the way users build and run applications. Technical and scientific Internet user communities as well as designers and developers of applications (especially Internet-based) can greatly benefit from the promise of a self-managed, self-scaling hosting environments.

As of today, the communities of multi-task computing, scalable communication systems, etc. have several well-known forums, but they focus mostly on system-driven scalability that could apply to specific scenarios such as datacenter companies, telecom operators, etc. Rather, easing developer's lives when transferring this scalability in a dynamic and pay-as-you-go manner to their application has received little (if any) attention.

## CFP (up to 300 words)

Please note that an Internet available demo or live video (hosted by the authors) would be very much appreciated and would very positively affect the review outcome.

This special issue aims to promote novel research on scalable Cloud Computing systems, especially in levels that have remained unexplored so far (beyond VM replication and/or automatic software reconfiguration). The list includes (but it is not restricted to):

1. Monitoring scalability for cloud infrastructure/platform/applications as compared to previous approaches.
2. Increased-scale data management solutions.
3. Improved resource allocation in the datacenter: resource leasing, statistical multiplexing, etc.

4. Dynamically scaling networks (e.g. exposing APIs for developers to demand, and pay, for more bandwidth at given points) and storage.
5. Scalable programming tools and models for cloud computing (e.g. increased scalability firewalls, etc.).
6. Application server scalability in PaaS cloud environments: multi-tenant containers, enhanced thread management services, etc.
7. Scalable cache systems.
8. Automated database replication and other highly scalable SQL and NoSQL forms.
9. Database isolation levels and their effects on performance and scalability.
10. Scalable support for ACID transactions and data replicas management.
11. Use cases of increased application scalability (before the cloud and after its migration to the cloud, using mechanisms other than VM replication).
12. Improved management tools to deal with cloud applications throughout their entire lifecycle (e.g. scaling decisions making modules or application governance modules that can be exposed as a service).
13. Languages and tools for expressing scalability.
14. Mobile application scalability.
15. Evolution of multitenant security as a function of the number of tenants (scale).

### **Instructions for Authors**

- The editors of the special issue are Rajkumar Buyya, Luis M. Vaquero, and Luis Roderó-Merino.
- Please submit paper to [Manuscript Central](#) as CloudScale special issue by December 15, 2011
- Notification of Acceptance/Reviewer comments will be given by March 31, 2012
- Electronic Publication: four months after all final papers submitted
- There is an 18 page length limit (12 point single space inclusive of figures and tables).
- Wiley has [Latex templates](#) but no special templates for Word; most papers are submitted in Word. Either Latex OR Word accepted.

### *Important Dates*

Dec 15, 2011 - Manuscript

Mar 31, 2012 - Notifications to Authors

Apr 30, 2012 - Revised Manuscript Submission

May 15, 2012 - Final Decision

Jun 15, 2012- Camera Ready Papers

### **Guest Editors**

Prof Rajkumar Buyya

University of Melbourne, Australia

Email: [raj@csse.unimelb.edu.au](mailto:raj@csse.unimelb.edu.au)

Website: <http://www.buyya.com/cv.html>

Dr. Luis M. Vaquero

Hewlett-Packard Labs.

Email: [luis.vaquero@hp.com](mailto:luis.vaquero@hp.com)

Website: <https://sites.google.com/site/luismvaquero/>

Prof. Luis Rodero-Merino

Universidad Politecnica de Madrid

Email: [lrodero@fi.upm.es](mailto:lrodero@fi.upm.es)

Website: <http://sites.google.com/site/luisroderomerinowebpage/research>