

GUEST EDITORIAL PREFACE

Special Issue on the 9th IEEE International Conference on Autonomic and Trusted Computing (ATC 2012)

Katinka Wolter, Fachbereich Mathematik und Informatik, AG Technische Informatik, Berlin, Germany

Philipp Reinecke, Fachbereich Mathematik und Informatik, AG Technische Informatik, Berlin, Germany

Vincenzo de Florio, Department of Mathematics and Computer Science, University of Antwerp, Antwerp, Belgium

This Special Issue collects extended versions of six papers presented at the 9th IEEE International Conference on Autonomic and Trusted Computing (ATC 2012) in Fukuoka, Japan, September 04-07, 2012. The papers have been carefully selected to provide an in-depth exploration of crucial issues in autonomic and trusted Cloud computing.

Cloud computing is characterised by highly-distributed computing and storage resources that are dynamically combined to achieve high flexibility in providing on-demand performance for large-scale applications. The incredible complexity of these systems cannot be handled by human administrators anymore and requires autonomous approaches to ensure dependability, performance, and security.

We set out with a study on software rejuvenation. Software ageing due to the accumulation of small errors affects many complex systems,

causing performance degradation and eventual failure. The high complexity of Cloud systems makes them especially susceptible to software ageing, as it is often impossible to eliminate all of the faults that lead to ageing. Software rejuvenation can help address these issues and ensure performance and dependability. In software rejuvenation, the system is restarted periodically. As restart resets the system to an error-free state, performance is restored and failure can be averted. On the other hand, the system is unavailable during the restart. This tradeoff prompts the question of how to find an optimal rejuvenation policy. In their paper, Hiroyuki Okamura, Tadashi Dohi and Kishor Trivedi optimise rejuvenation policies with respect to interval reliability, encompassing previous measures such as the point-wise system availability.

The second paper addresses the question of resource allocations in the Cloud. In Cloud

computing, customers demand computing and storage resources to ensure high performance of their tasks; providers, on the other hand, need to reduce costs by consolidating resources and eliminating idle times. In order to achieve optimal performance and minimal costs, resource allocations must thus be carefully optimised. Wei Chen, Xiaoqiang Qiao, Jun Wei, Hua Zhong and Tao Huang study optimisation of resource allocations in the Cloud and provide a method for achieving optimal static allocations and for reducing the costs of dynamic reconfiguration.

Dynamic resource allocation is one of the most important aspects of Cloud systems, as it allows applications to react to fluctuations in the workload by rapidly allocating resources to maintain performance as the load increases, or releasing resources when the load decreases, thus reducing costs and energy consumption. Apart from providing performance, dynamic scaling of resources can also affect availability of the application. Wenting Wang, Heopeng Chen and Xi Chen address this area with a new availability-focused approach for scaling resource allocations.

The management of resource allocations in the Cloud requires accurate low-overhead monitoring of the performance of virtual machines. Performance degradation due to overload, in particular, must be detected as quickly as possible, in order to be able to scale up resource allocations in time to minimise the impact on the quality of service for the customer. In the fourth paper, Toshiaki Hayashi and Satoru Ohta present an approach to infer degraded performance of a virtualised server from passive measurements of the network traffic of the virtual machine. This approach enables detection of performance degradation even in high-load situations where monitoring solutions inside of the virtual machine may fail to report accurate results.

Cloud computing provides an important basis for many complex, resource-hungry scientific tasks. In order to automate these tasks,

they are typically organised as workflows, which are then executed by a workflow execution engine that leverages properties of the underlying platform such as parallelism for optimising performance and dependability. The paper by Luis Assuncao, Carlos Goncalves and Jose C. Cunha proposes a new framework for workflow management and illustrates its application on a real-world Cloud service.

Security has become an ever more important issue in recent years. In Cloud systems, security breaches can have extremely expensive consequences, as they affect a large number of customers and can undermine end-users' trust in such systems, eventually leading to economical failure. Furthermore, the vast resources of Cloud systems make them particularly attractive as stepping-stones in attacks on other systems. The final paper in this volume presents an approach for detecting, locating and mitigating stepping-stone attacks. Kenichi Kourai, Takeshi Azumi, and Shigeru Chiba describe a method that combines packet-filtering in the virtual machine monitor with virtual-machine introspection for identification of the originators of malicious data.

This special issue followed a conference with more than 50 submissions, from which 11 extended papers were considered. The present selection of papers was determined after a rigorous reviewing process, and we wish to express our sincerest gratitude to the reviewers for their great work. We hope this volume inspires the reader regarding the challenges of the need for autonomous and secure solutions for Cloud computing, and we will be delighted if it fosters further development in this exciting area.

Katinka Wolter
Philipp Reinecke
Guest Editors
Vincenzo de Florio
Editor-in-Chief
IJARAS