

Methodological Principles for Reproducible Performance Evaluation in Cloud Computing

Alessandro V. Papadopoulos
Mälardalen University, Sweden

Laurens Versluis
Vrije Universiteit Amsterdam,
The Netherlands

André Bauer
Nikolas Herbst
Jóakim von Kistowski
University of Würzburg, Germany

Ahmed Ali-Eldin
Umeå University, Sweden

Cristina L. Abad
Escuela Superior Politecnica del
Litoral, Ecuador

José Nelson Amaral
University of Alberta, Canada

Petr Tůma
Charles University, Czech Republic

Alexandru Iosup
Vrije Universiteit Amsterdam,
The Netherlands




Figure 1: Presence of selected principles in the cloud performance evaluation.

ABSTRACT

The rapid adoption and the diversification of cloud computing technology exacerbate the importance of a sound experimental methodology for this domain. It has been a decade since the first commercial cloud has opened for a general public (Amazon AWS, in 2007). Cloud computing systems are now in much demand, and command significant industrial and scientific interest. Although scientific progress and industry growth depend on using sound principles for measuring and reporting cloud-system performance, this process remains complex. This work investigates how to measure and report performance in the cloud, and how well the cloud research community is already doing it. There are several difficulties, including controlling the time-accuracy trade-off, deciding which data to collect and to report, and using the right summarization statistics in the report. We propose a set of eight important methodological principles that combine best-practices from nearby fields with concepts applicable only to clouds, and with new ideas about the time-accuracy trade-off. We show how these principles

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICSE '20, May 23–29, 2020, Seoul, South Korea

© 2020 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00
<https://doi.org/10.1145/1122445.1122456>

are applicable using a practical use-case experiment. To this end, we analyze the ability of the newly released SPEC Cloud IaaS benchmark to follow the principles, and showcase real-world experimental studies in common cloud environments that meet the principles. Last, we report on a systematic literature review including top conferences and journals in the field, from 2012 to 2017, analyzing if the practice of reporting cloud performance measurements follows the proposed eight principles. Worryingly, this systematic survey and the subsequent two-round human reviews, reveal that few of the published studies follow the eight experimental principles. We conclude that, although these important principles are simple and basic, the cloud community is yet to adopt them broadly to deliver sound measurement of cloud environments.

ACM Reference Format:

Alessandro V. Papadopoulos, Laurens Versluis, André Bauer, Nikolas Herbst, Jóakim von Kistowski, Ahmed Ali-Eldin, Cristina L. Abad, José Nelson Amaral, Petr Tůma, and Alexandru Iosup. 2020. Methodological Principles for Reproducible Performance Evaluation in Cloud Computing. In *ICSE '20: ACM International Conference on Software Engineering*, 23–29, 2020, Seoul, South Korea. ACM, New York, NY, USA, 1 page. <https://doi.org/10.1145/1122445.1122456>

REFERENCES

- [1] A. V. Papadopoulos, L. Versluis, A. Bauer, N. Herbst, J. Von Kistowski, A. Ali-eldin, C. Abad, J. N. Amaral, P. T. Árma, and A. Iosup. 2019. Methodological Principles for Reproducible Performance Evaluation in Cloud Computing. *IEEE Transactions on Software Engineering* (2019), 1–1. <https://doi.org/10.1109/TSE.2019.2927908>