

Benchmarking the Cloud in a simple and effective way

Benchmarking is a well-known and extensively used technique for assessing and comparing performance of systems. Applying benchmarking to Cloud services creates new opportunities to make a more efficient usage of the Cloud and save costs. In fact, performance of Cloud resources heavily affects the performance of the applications running in the Cloud. Knowing the resources performance is a factor to take into consideration to estimate the costs to sustain in order to guarantee the desired application Quality of Service.

Performance can sensibly vary from Cloud to Cloud and it is difficult to make comparisons on paper because actual performance depends on factors like the hardware equipment of the Cloud infrastructure, resources co-location and the configuration of the assigned resources. Unfortunately, most of this information is not disclosed by Cloud Providers. For this reason, running independent benchmarking tests remains the most reliable way of comparing performance of different Cloud Providers.

An extensive and complete benchmarking of Cloud poses multiple challenges. First of all, it is necessary to generate synthetic workloads that resemble the application behaviour (because performance can really vary from workload to workload). Then, the benchmarks need to be executed repeatedly in different Clouds and their results need to be aggregated and compared. When taking into account the multitude of Cloud Providers, in conjunction with the different resource types and sizes that are offered and that need to be analysed, it becomes clear that the benchmarking of Cloud Providers is a time consuming and non trivial task.

The Benchmarking Suite, developed by [Engineering Ingegneria Informatica S.p.A.](#) in the context of the two EU research projects [ARTIST](#) and [CloudPerfect](#), aims at addressing these challenges providing an all-in-one solution for the benchmarking of Cloud resources simulating common application behaviours and comparing the observer performance. The Benchmarking Suite embeds a set of authoritative, de-facto standard and widely used third-party benchmarking tools and relies on them for the workload generation and the performance measurement. The Benchmarking Suite automates the benchmarking process, managing the allocation and de-allocation of necessary resources on the Cloud, the installation and execution of the benchmarking tools and the collection of result data.

The Benchmarking Suite handles multiple tools, multiple Cloud Providers and multiple benchmarking sessions at the same time managing the specificities of each of them. It provides the user with a unique, coherent and simplified interface to execute any test on any Cloud. It has been designed to be extensible and to allow an easy integration of further third-party benchmarking tools and Cloud Providers. Data collected and stored during the benchmarks execution is homogenized and aggregated

on different higher-level metrics (e.g. average values) allowing performance comparison among different Clouds.

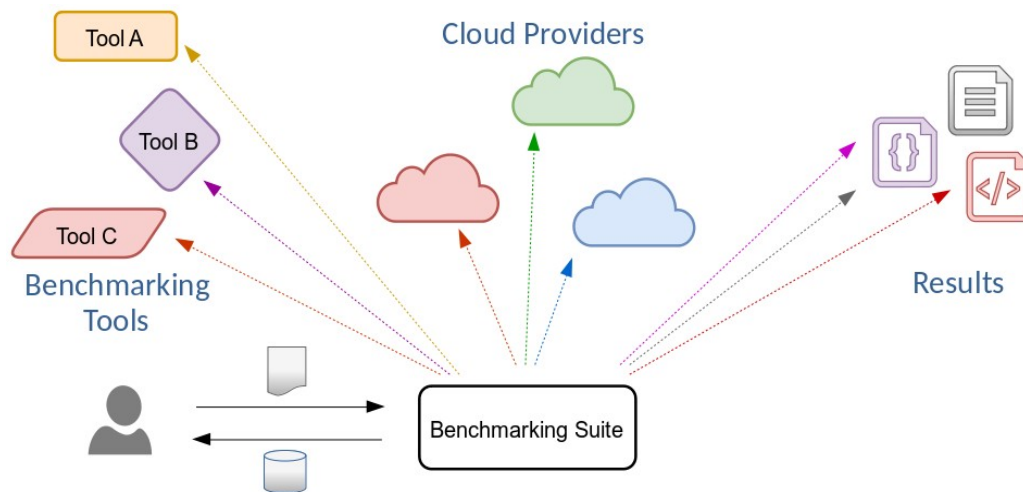


Illustration 1: The Benchmarking Suite is a single access point to Cloud benchmarking that automates and simplifies the execution of tests and the analysis of results.

The execution of the same benchmarks on different combinations of Clouds, resource types and time periods is a common task when the goal is to find the best offering for a specific application. The Benchmarking Suite supports this task through the [Benchmarking Scheduler](#) that allows to select and configure in advance the benchmarks and to schedule their execution on a regular basis.

The Benchmarking Suite is being used in the the EU Innovation Action [CloudPerfect](#) to assess public and private IaaS Clouds. The performance data generated by the Benchmarking Suite is analysed and aggregated to compute higher-level metrics (e.g. Virtual Core Performance) and to provide customized rankings of Clouds on the basis of the application nature.

The Benchmarking Suite is a pure Python application and uses some third-party libraries like [Apache LibCloud](#) to access cloud services and [Paramiko](#) to execute commands on remote Virtual Machines. It offers two different user interfaces: a command-line tool and a REST API. It is distributed through [PYPI](#) and [Docker Hub](#).

References

- Source Code: <https://github.com/benchmarking-suite>
- Documentation: <http://benchmarking-suite.readthedocs.io/>
- Support: <https://github.com/benchmarking-suite/benchsuite-issues/issues/>

Legal

The Benchmarking Suite is released under the [Apache 2.0](#) license.



Author Info

Gabriele Giammatteo, [M] graduated in 2009 in Computer Engineering at the “Università degli studi di Roma Tre” on secure storages in distributed environments. He is employed in Engineering's R&D Lab as software engineer, working in the setup and management of distributed infrastructures in particular in cloud computational and data management systems. He has been involved in EU FP7 and H2020 projects since 2009 focused on various topics of cloud computing software engineering ranging from software integration, testing and quality assurance (D4Science, iMarine, OCEAN, S-CASE, BlueBRIDGE), to benchmarking (ARTIST, CLOUDPERFECT), to automation (TEFIS, EGI-ENGAGE), to cloud ecosystem sustainability (OCEAN).