PRESEnCE: Performance Evaluation of SaaS Web Services across Multi-Cloud Providers

Abdallah Ali Z. A. IBRAHIM, Sebastien VARRETTE, Dzmitry KLIAZOVICH and Pascal BOUVRY

Abstract

Cloud services providers deliver cloud services to cloud customers on a pay-per-use model while the quality of the provided services are defined using Service Level Agreements also known as SLAs. Unfortunately, there is no standard mechanism which exists to verify and assure that delivered services satisfy the signed SLA agreement in an automatic way, which impedes the possibility to measure accurately the Quality of Services (QoS). In this doctoral thesis, we aim at offering an automatic framework able to evaluate the QoS and SLA compliance of Web Services offered across several Cloud providers.
The ambitions are three-fold:

1. Quantify in a fair and by stealth way the performance and scalability of the delivered Web Services. By stealthyness, we refer to the capacity of evaluating a given service through multiple workload patterns that makes them undistinguishable from regular traffic for the provider point of view.
2. Assess the claimed SLA and the corresponding QoS from a set of relevant performance metrics, in particular to verify the service response time.
3. Provide a multi-objective analysis of the gathered performance metrics to be able to classify cloud brokers and help an end-user to select the appropriate providers based on valuable and measurable criteria.

In this context, the PRESEnCE framework (PeRformance Evaluation of SErvices on the Cloud) is proposed.
It defines a set of common metrics for measuring the behavior of services/applications while being delivered by specific cloud providers. Those metrics have been designed to assess the providers' scalability and performance for the listed services. In practice, several reference benchmarks have been considered and evaluated to characterize this behavior. So far, the following benchmarking tools have been integrated: Iperf, Twitter RPC-perf, Yahoo Streaming Benchmark and Apache HTTP server benchmarking tool. Because of the difficulty of assuring and verify the SLA compliance between cloud providers, one of the challenges of this PhD is to find a way to combine these tools by stealth to evaluate cloud application performance. Indeed, recent scandals in the automotive sector demonstrate the capacity of solution providers to adapt the behavior of their product when submitted to an evaluation campaign. We want to prevent such actions in the Cloud sector by rendering the performance evaluation of SaaS-based services as close as possible from a regular yet heavy usage of the considered service. Benefits are foreseen for both the application providers (to assess the performance and SLA appliance of their offer or the one of their competitors), or for the end-users as the PRESEnCE framework permits to control and validate the SLA and corresponding QoS.

In order to achieve the afore-mentioned goals, this doctoral thesis requires an in-depth understanding of "regular" Cloud usage patterns and ways to measure the performance of given web services. Our first results in these domains are presented in this study. In particular, details of the deployed testbed, together with the evaluation component of the PRESEnCE framework are proposed.