

# Industrial workload replay for FaaS platforms

**Keywords:** cloud, Function-as-a-Service, performance analysis, Docker, statistics

## Context

Cloud, i.e., renting remote computing resources, is the main method for deploying applications. It has eventually reflected on the *architecture* of the applications themselves: we observe a trend of “cloud native” designs. The newest paradigm for cloud applications is *Function-as-a-Service* (FaaS): features are served by composing and replicating instances of atomic functions.

Research data, i.e., real workload traces, remain scarce; and for the rare published ones, the researchers lack *performance analysis tools* that could help evaluating novel work on these data.

## Goal

FaaSLoad is a research software to inject a workload and monitor the performance in a FaaS platform. It is already in working shape, including its main features: injecting a workload, building a dataset of function executions, and monitoring performance and resource usage.

Nonetheless, despite FaaSLoad’s ability to play compatible workload traces, it cannot *replay existing workload traces* yet. The challenge lies in the nature of the data: instead of precise and exhaustive traces providing complete events, the traces’ publisher (Microsoft Azure) “summarized” them into *statistical properties* (such as invocation frequency, average resource usage, etc.) over time. Moreover, the actual workload – the executed cloud functions, are of course not provided, for confidentiality reasons.

The goal is thus to develop the ability to replay those traces. More precisely:

- to develop a dummy cloud function to fake the resource usage;
  - some previous work done;
- to study the industrial workload traces and to understand their content and structure;
  - some previous work done;
- to develop a converter for industrial workload traces, to produce traces that FaaSLoad can replay;
- to ensure the produced traces conform to the statistical properties of the original industrial traces;
- to ensure the actual replayed workload conforms to the statistical properties of the original industrial traces.

The project is to be supervised by Mathieu Bacou ([mathieu.bacou@telecom-sudparis.eu](mailto:mathieu.bacou@telecom-sudparis.eu)).