

# Experiences of Cloud Storage Service Monitoring: Performance Assessment and Comparison

Enrico Bocchi<sup>1,2</sup>, Idilio Drago<sup>1</sup>, Marco Mellia<sup>1</sup>

*1 - Politecnico di Torino, Italy*

*2 - Telecom ParisTech, France*

{name.surname}@polito.it

---

## Abstract

Personal cloud storage services offer a large amount of space and the ease to synchronize devices with great simplicity. They help in sharing content and in backing up data by relying on the cloud to seamlessly store users' files. Despite the high public interest in such services, little information is available about design, implementation choices and, most of all, actual performance implications.

In this work, we design a methodology to run benchmarks of cloud storage services. We unveil how they are implemented, where their servers are located, and measure implication on performance as seen by end-users. By means of repeatable and customizable tests, we identify eventual advanced capabilities the cloud client implements, and measure performance implications. We consider realistic workloads (e.g., the exchange of text and binary files, compressed archives, the presence of file replicas, etc.) and network accesses (e.g., high speed university campus, or 3G mobile connectivity).

We use then the benchmarking methodology to compare 11 cloud services, including popular solutions like Dropbox or Google Drive, and two private storage solutions, i.e., the open source ownCloud and the commercial VMware Horizon, that we installed and configured in our campus network. We take the perspective of a customer located in Europe, and we benchmark each service.

Our case study reveals interesting differences in design choices. Results show no clear winner, with all services having potential for performance improvements. Some are limited by design choices, e.g., by artificially throttling upload and download speed, or by long application timers that slow down synchronization procedures. Others suffer TCP performance issues due to their data centers being located in other continents. In some scenarios, the synchronization of the same set of files can take 20 times longer. In other cases, we observe a wastage of twice as much network capacity, questioning the design of some services, especially in a bandwidth constrained scenario like 3G/4G connectivity.

Our results show the implications of design choices on performance, and of the trade-offs faced when building cloud storage services. The developed methodology and the collected results are useful both as benchmarks and as guidelines for system design. In addition, they help the prospected customer in the choice of the best service by allowing the execution of independent performance tests before purchasing a storage offer.

*Keywords:* Cloud Storage; Dropbox; Measurements; Performance; Benchmarking

---