

Master 2 Internship Subject

Cost-Aware Replication of Distributed Data for Cloud Computing

Description

In many distributed systems, data are replicated for the sake of performance, reliability, availability and the like. For a given application and a given distributed infrastructure, the number of replicas of any data item might be different, and the optimal location of the replica differ with the context of use and time. Indeed, replica might move from site to site for instance to follow the associated computation, or the number of replica might change dynamically to adjust to the load and reliability of the distributed system. Associated to the replica and their management is the notion of costs, embedding for instance the cost to host the data items in a Cloud storage system. These costs reflect the use of the different resources needed to handle a given replica, or a set of replicas. They can be interpreted in \$ or €, accounting most of the time for the amount of disk resource used and the network bandwidth associated. On the other hand, replicas also increase the quality of service for the cloud tenants, and results in less penalties paid by the cloud providers regarding the SLA agreed with its tenants. Given a set of context information (cost of hosting, networking, number of sites, penalties) it is possible to derive an optimal number and location of the replicas.

In the context of this work, we plan to add the power and energy needed, in the form of constraints (i.e. the total power used by the replicas might not exceed a given electrical budget, which is called power-capping, or the total energy for storing the data for a long time might not exceed an energy budget, which is called energy-capping), or in the form of objectives (i.e. minimising the energy used without loss of performance, or under a performance constraint). Since electricity or energy can easily be translated to \$ or €, depending on the location, the season and even the time of the day, including these fluctuations in the planning of replica management can be done. Then, power and energy consumptions must be included in the provider expenditures that are modelled in the economical cost model. The challenge consists in (i) understanding how the power or energy constraints adds to the difficulty of the planning and change the decision processes, (ii) understanding their impact on the long term for the cloud provider and its tenants and (iii) considering the trade-off between the tenant objective satisfaction and the provider profit meeting while taking into account these new constraints.

Keywords

Cloud Computing, Data Partitioning, Data Replication, Energy Consumption, Economical Cost Model, Trade-off modeling.

Supervisors

Jean-Marc Pierson and Riad Mokadem (Jean-Marc.Pierson@irit.fr, Riad.Mokadem@irit.fr)

Internship Start date and duration

March 1st, 2018. 5 months.

Possibility of continuation in a PhD.