

Cloud Computing Resilience in practice: harder than it looks

Paulo Veríssimo, Alysson Bessani, Marcelo Pasin
University of Lisbon Faculty of Sciences
Lisbon, Portugal
{pju,bessani,pasin}@di.fc.ul.pt

Cloud Computing (CC) has gained enormous momentum and is progressively reaching every sector using IT. As more and more services migrate to CC, so increases the dependence of the IT business on the latter, perhaps not yet met by adequate levels of robustness. This can be testified by the numerous failures of cloud provider services made public, having caused service and data loss, as well as confidentiality compromises. Nevertheless, short of promising adequate security management of the infrastructure and perhaps some form of disaster recovery, little more has been offered by cloud providers so far.

How is this being addressed today in deployed systems?

(i) Approaches confined to single cloud provision will not address high-resilience objectives, since they are a single point of failure. Large cloud providers attempt to achieve resilience by deploying several, differently located, cloud subsystems (i.e., in different datacenters), which definitely improves on the situation, but they still remain under a single management and trust domain, with regard to common-mode and malicious faults, which may result, at organizational level, in single points of failure (e.g., recent Amazon and Azure outages affected more than one availability zone).

(ii) Proprietary trusted or accredited clouds may implement specific IaaS or PaaS approaches to achieving resilience. However, they can make migration or interoperation difficult and expensive, creating vendor lock-in and competition exclusion.

(iii) Federated cloud environments certainly introduce the next instance of a solution, but they require alliance of the involved providers, potentially reiterating the problem of common or related management and trust domains. Also, the general reality is that providers usually compete and may be mutually distrusting or even hostile to each other.

Worthwhile mentioning are recent research efforts on the *cloud-of-clouds* computing paradigm as a path to achieve cloud computing resilience. The cloud-of-clouds paradigm extends the cloud concept, by leveraging the availability of multiple independent cloud environments to create diverse ecosystems, and by letting users (besides providers) self-organize the way they use multiple cloud computing offerings.

For the above reasons, alternative avenues are desired, and such a discussion is timely. Interesting topics to discuss in this session will be those addressing the problems above, for example:

- What are the adequate paradigms to achieve cloud computing resilience?
- How to design built-in, open, and diverse solutions to cloud dependability and security that preserve legacy?
- How to achieve resilience against both attacks and accidents, doing so in as automated as possible a way, but keeping performance at acceptable levels?