# Disaster-Resilient IaaS Cloud Networks: Requirements and Research Directions

**Rodrigo S. Couto**[1], **Stefano Secci**[2], **Miguel Elias M. Campista**[1] **e Luís Henrique M. K. Costa**[1] *

[1]Universidade Federal do Rio de Janeiro - PEE/COPPE/GTA - DEL/POLI
[2]Sorbonne Universités, UPMC Univ Paris 06, UMR 7606, LIP6, F-75005, Paris, France

{souza,miguel,luish}@gta.ufrj.br, stefano.secci@lip6.fr

***Abstract.*** *Many corporations are migrating their IT infrastructure to the cloud by using IaaS (Infrastructure as a Service) services. Nevertheless, with IaaS, the company relinquishes control of the physical infrastructure. Therefore, it can only rely on IaaS services if providers can guarantee performance and security levels. To encourage IaaS subscriptions, cloud providers must employ a resilient data center. To this end, IaaS providers deploy a lot of redundancy on their infrastructure, to overcome various types of failures, such as hardware (e.g., failure in hard disks, network cables, and cooling systems), software (e.g., programming errors), and technical staff (e.g., execution of wrong maintenance procedures). This strategy, however, does not guarantee service availability under force majeure and disaster events that are out of the provider's control. Although IaaS providers often do not consider catastrophic events, they can offer recovery services such as virtual machine replication and redundant network components to reduce the risk of data loss and virtual machine unavailability. These services can be provided as long as a data center infrastructure resilient to disasters is available, which is generally composed of several sites spread over a region. The resilience of multiple data center sites depends on the underlying network infrastructure employed to interconnect them. The literature in this domain is still scarce, although it is a fundamental problem to allow business continuity for IaaS clients. To motivate work in this area, in this work we propose guidelines to design a disaster-resilient cloud network to support IaaS services. These guidelines describe design requirements, such as the amount of data loss tolerated upon a disaster, and allow the identification of research directions, such as disaster-resilient placement of physical servers and virtual machines over a wide area network.*

*This work was partially supported by FAPERJ, CNPq, CAPES research agencies.