

Themis

A Spot Market-based Automatic Resource Scaling Framework

Stefania Costache, Nikos Parlavantzas, Christine Morin, Samuel Kortas

Inria Rennes - Bretagne Atlantique/IRISA - Myriads team, EDF R&D, France

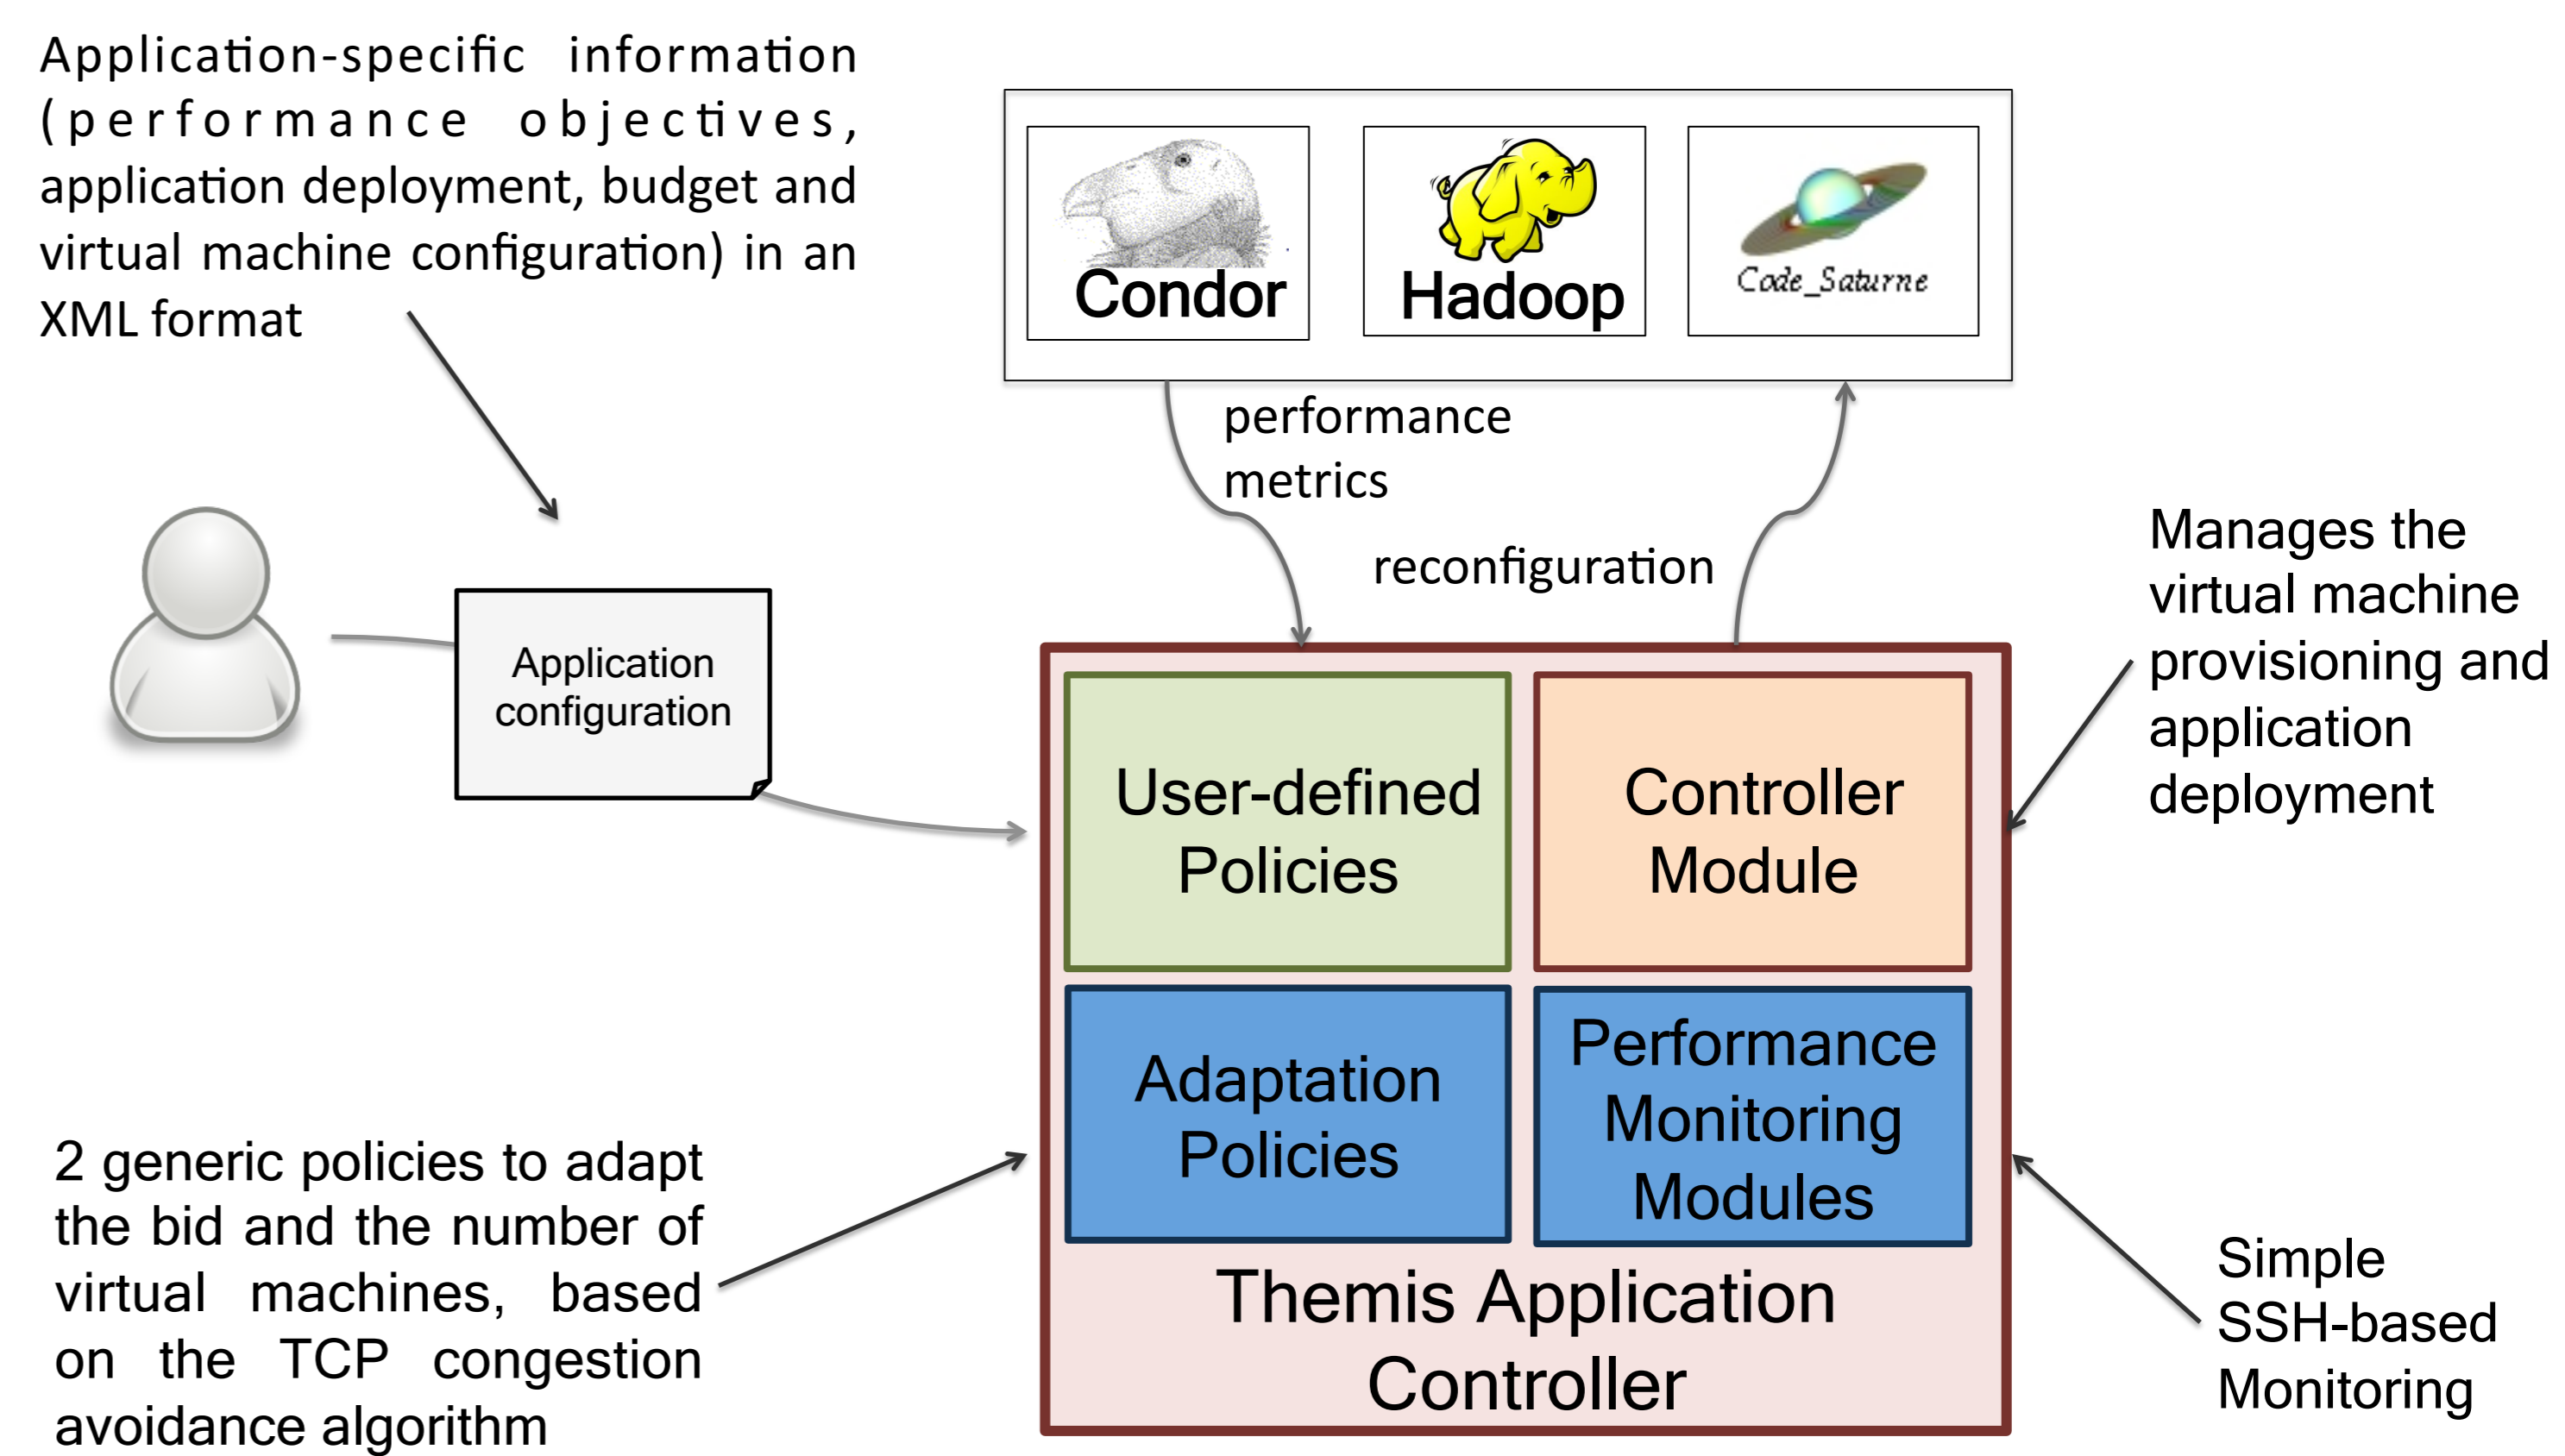
Problem

- Scientific applications are becoming more complex and dynamic.
- To support these applications resource management mechanisms are required to be **simple and flexible**.
- Current mechanisms used by cloud providers meet these requirements as they allow users to provision virtualized resources **instantly**. However they **DON'T**:
 - provide support for building dynamic applications and meeting user SLOs
 - lead to an efficient resource usage.

Our Goals

- Provide generic support to applications to scale automatically their resource demand on cloud systems
- Enforce fair resource utilization among users
- Maximize the infrastructure utilization of cloud providers

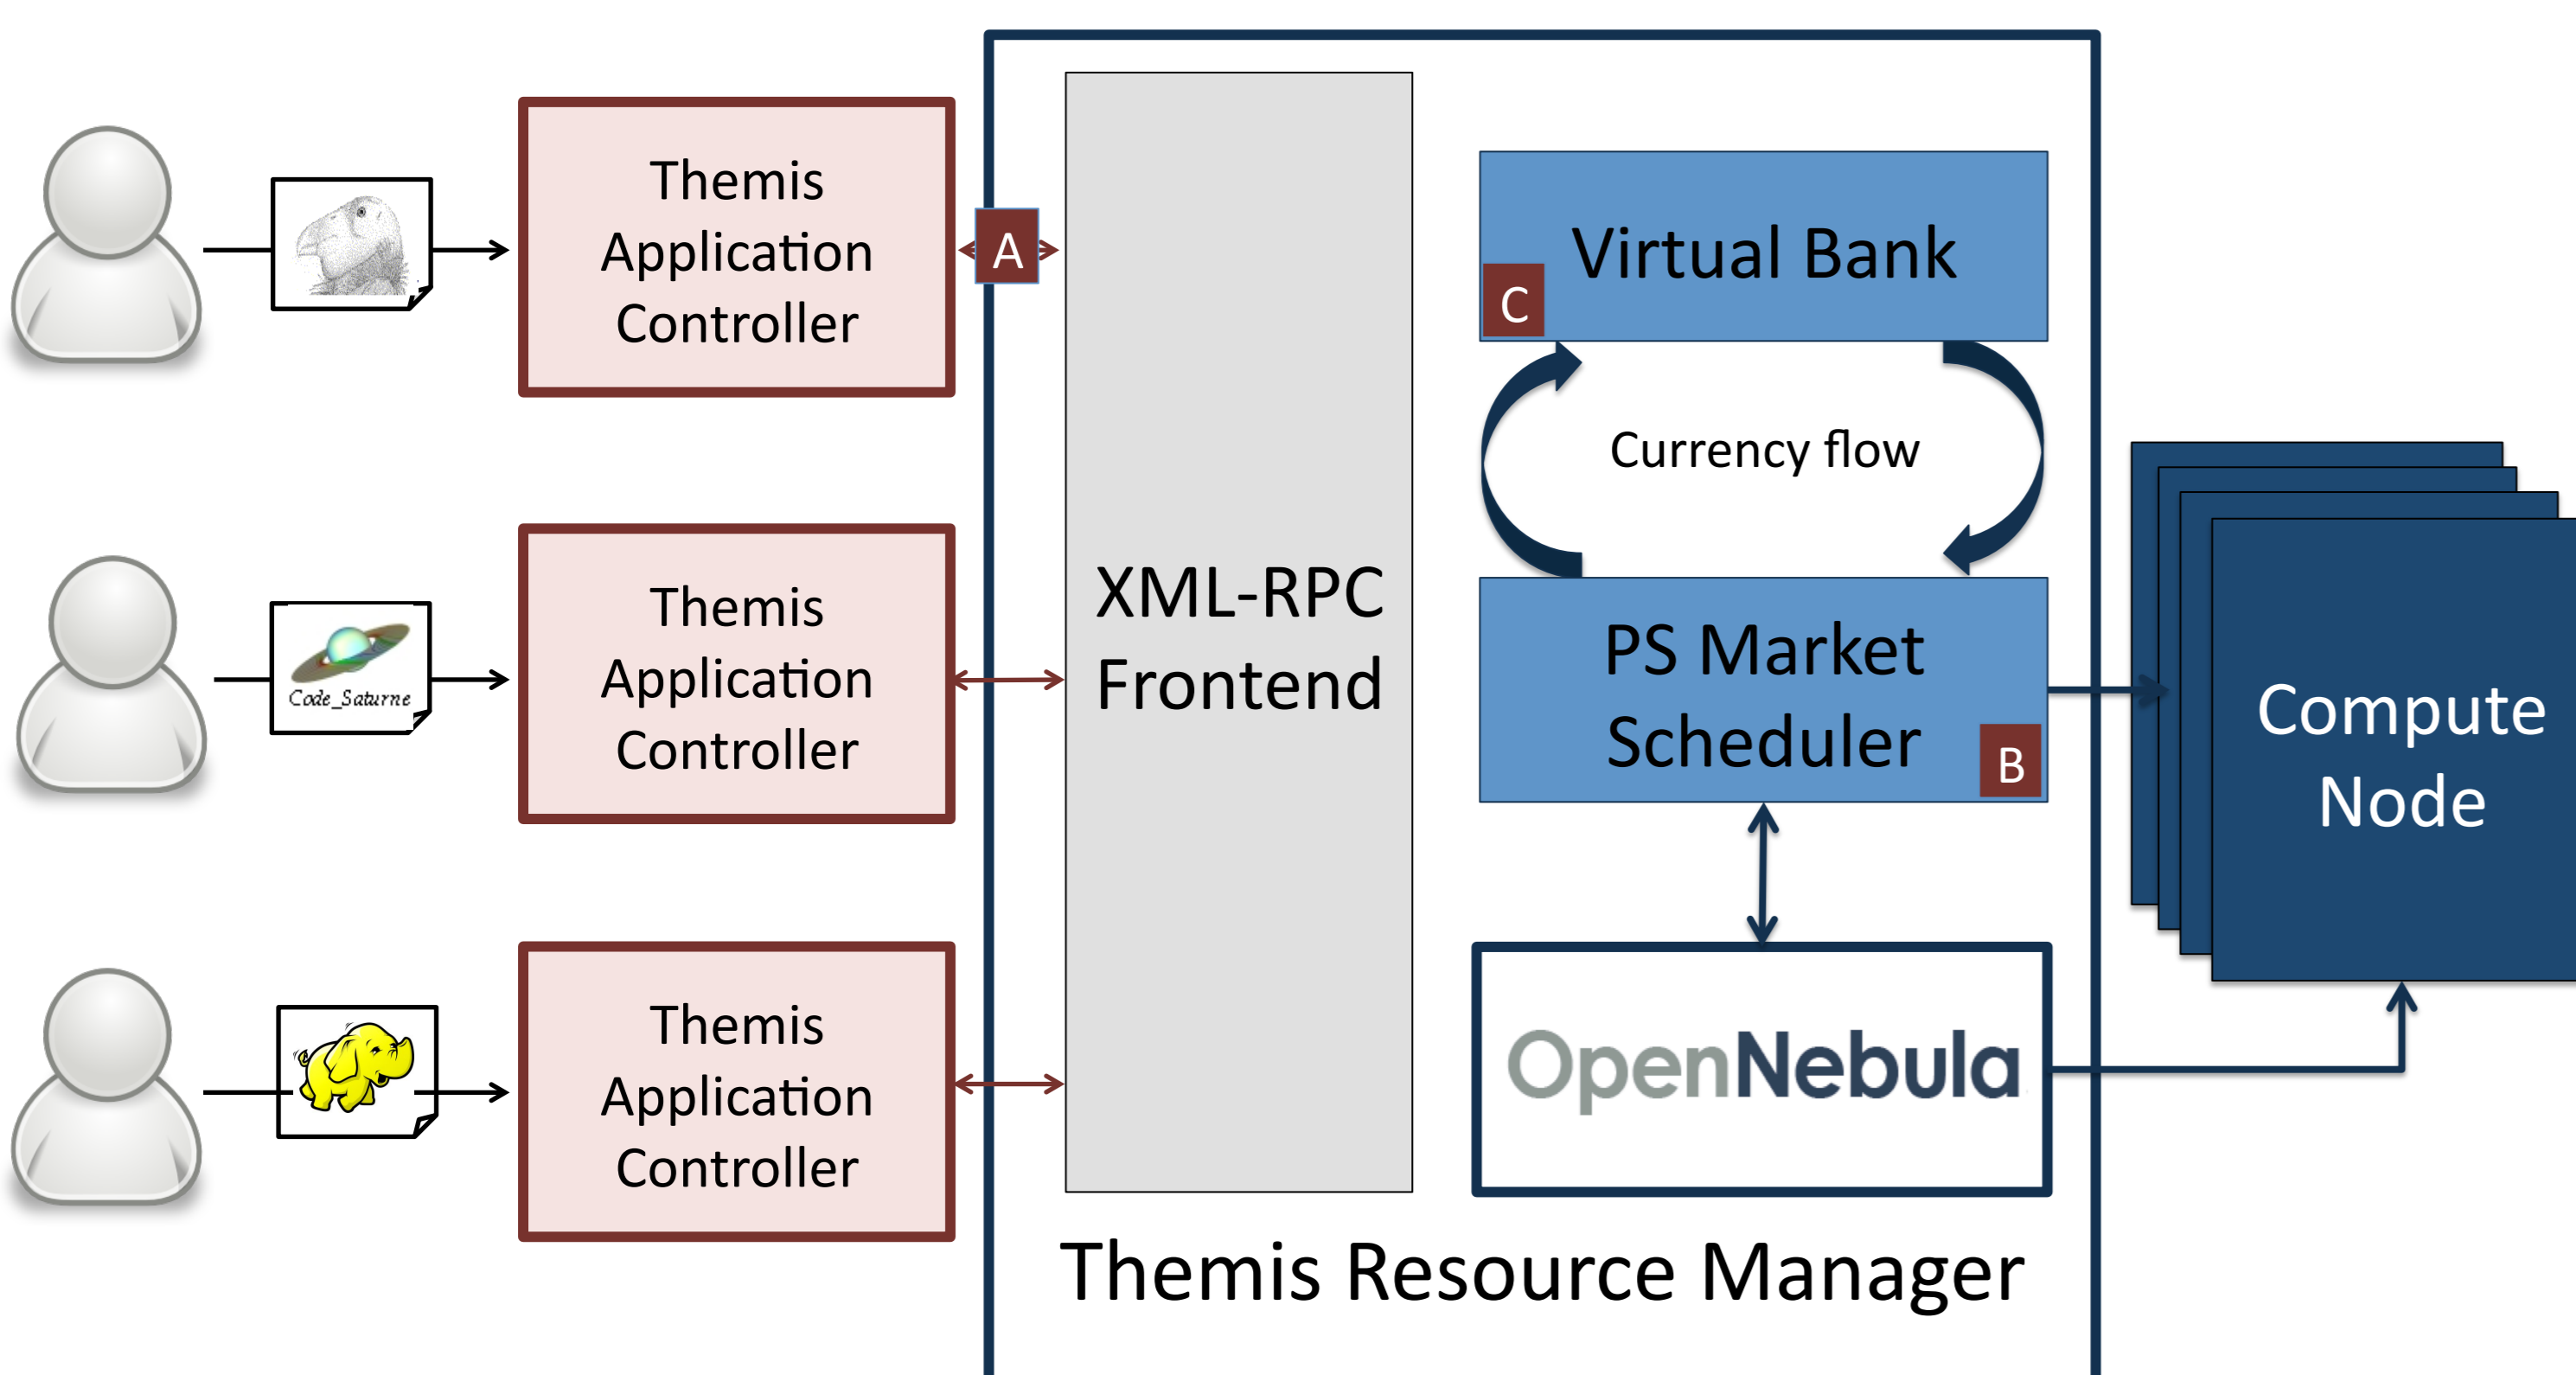
Autonomous Application Controllers



Themis

Themis is a decentralized resource control system for virtualized infrastructures based on:

- a **virtual economy** and a **proportional share (PS)**-based spot market to allocate virtualized resources (VMs), providing fairness to users and leading to a maximum resource utilization
- autonomous controllers** to scale the application resource demand to meet user SLOs




- A** Application controllers receive information about the current resource price and capacity and adapt the VM number and their bids to meet user SLOs.
- B** The PS Market Scheduler computes the VM allocations and the physical placement.
- C** The virtual bank enforces virtual currency management policies to avoid hoarding and starvation.

Resource Allocation through a Spot Market


- Users bid for VMs while the physical resource share allocated to each VM is proportional to the cloud resource price.
- The proportional-share market ensures a maximum resource utilization as it is **work-conserving** and allocates resources in a **fine-grained** manner.
- A load balancing algorithm ensures that the virtual machines get the best shares for which the users are willing to pay, while keeping a low number of virtual machine migrations.

Running an Elastic Application with Themis


We have set up a Condor pool and used a Themis controller to provision VMs according to the fluctuations in the number of jobs and the resource price.



The controller adjusts the number of VMs according to the number of running and waiting jobs.



The controller can buy a maximum number of VMs, according to the infrastructure capacity. In this case the Condor pool uses the entire infrastructure.



The controller adjusts the number of VMs according to the fluctuating price. When the price increases the controller provisions less VMs.

Acknowledgments

This work is supported by ANRT through the CIFRE sponsorship No. 0332/2010

stefania.costache@inria.fr nikos.parlavantzas@inria.fr christine.morin@inria.fr samuel.kortas@edf.fr