



HAL
open science

Grid'5000: A Production-grade Testbed for Experiment-driven Computer Science on HPC and Clouds

Lucas Nussbaum, Pierre Neyron, Olivier Richard, Emmanuel Jeanvoine

► **To cite this version:**

Lucas Nussbaum, Pierre Neyron, Olivier Richard, Emmanuel Jeanvoine. Grid'5000: A Production-grade Testbed for Experiment-driven Computer Science on HPC and Clouds. Inria Booth at SC'13, Nov 2013, Denver, United States. 2013. hal-00920389

HAL Id: hal-00920389

<https://inria.hal.science/hal-00920389v1>

Submitted on 18 Dec 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Overall Goal

Design & provide a testbed that supports high-quality experimental science

- ▶ Large scale, complex experiments
- ▶ Trustworthy results
- ▶ Reproducible experiments

1200 nodes (8000 cores) in 10 sites
500+ users per year since 2006

Key features

- ▶ Experiment on all layers : fully reconfigurable

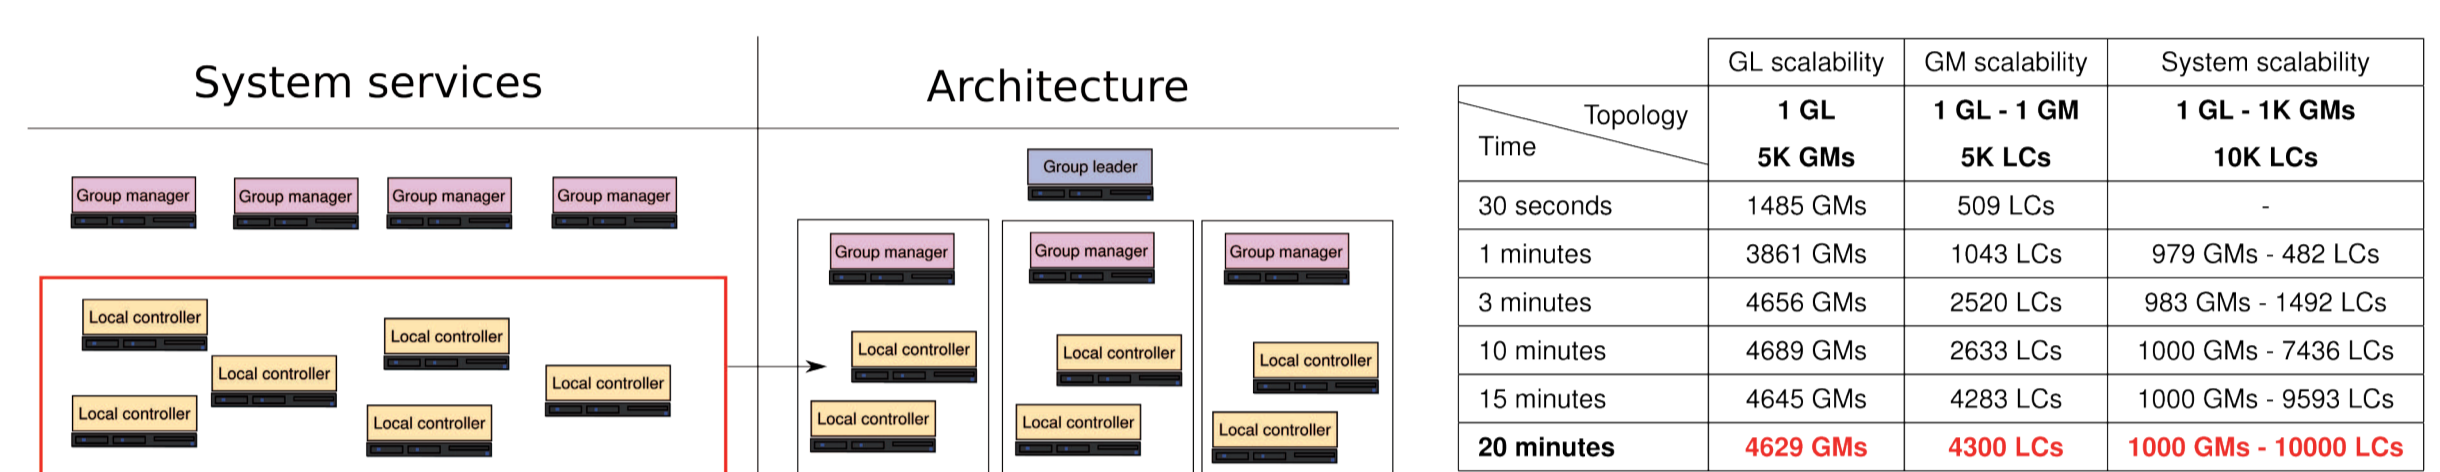


- ▶ **Software stack** (root on bare metal)
- ▶ **Network isolation** (intrusive security-sensitive protocols)
- ▶ **Experimental conditions** (fault injection)
- ▶ **Resources reservation** (nodes, network, storage)
- ▶ **REST API** for XP control tools (Expo, XPFlow, Execo)
- ▶ **Measurement and monitoring** infrastructure
- ▶ Support for **Big Data** and **Open Science**

- ▶ **Access high end hardware networks**
 - ▶ **High performance networks** : Infiniband, Myrinet
 - ▶ **Accelerators** (GPU, MIC)
 - ▶ **Dedicated 10 Gbps backbone** (RENATER)


Ground-breaking experiments

- ▶ **Snooze Autonomic Cloud Management System**



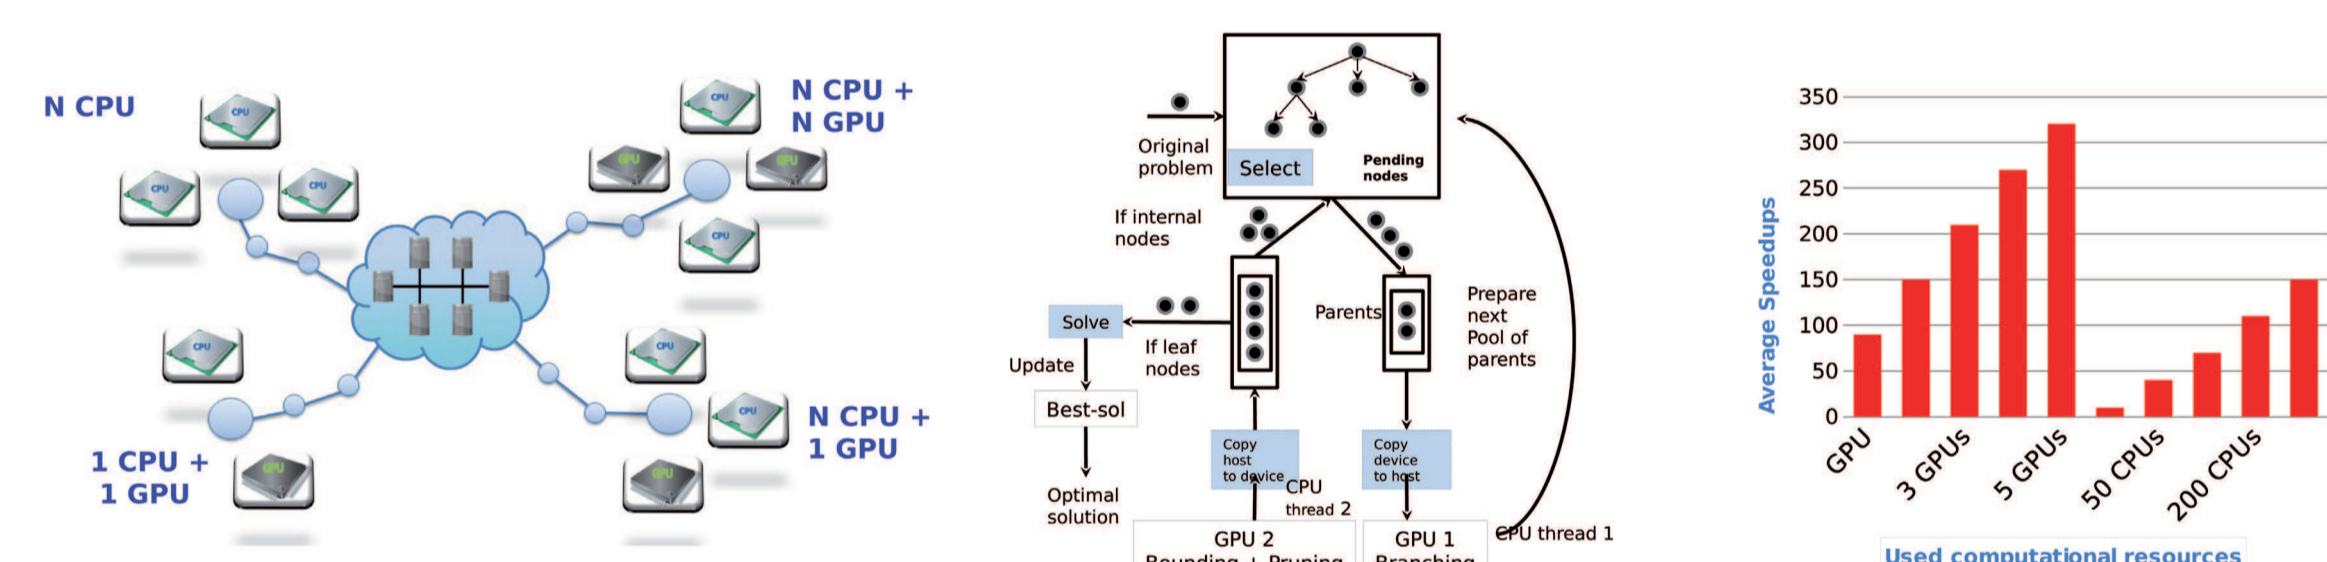
Scalability evaluation on 538 nodes / 5 sites

- ▶ **Testing of the Kadeploy3 provisioning system**




4000 virtual machines in a global VLAN

- ▶ **Solving combinatorial optimization problems**



Hybrid multi-core and GPU computing

- ▶ **Calibration of the SMPI simulator models**



Combining real and simulated executions

Get a free account

<http://www.grid5000.fr/open-access>

Software stack

- ▶ **Kadeploy** : <http://kadeploy3.gforge.inria.fr>
- ▶ **OAR** : <http://oar.imag.fr>
- ▶ **Distem** : <http://distem.gforge.inria.fr>
- ▶ **Expo** : <http://expo.imag.fr>
- ▶ **XPFlow** : <http://xpflow.gforge.inria.fr>
- ▶ **Execo** : <http://execo.gforge.inria.fr>