

# Deployment of Cloud stacks on Grid'5000

Lucas Nussbaum  
lucas.nussbaum@loria.fr



# Motivations

- ▶ Grid'5000:
  - ◆ 502 active users over the last year, 2762 overall
  - ◆ Mostly Computer Science researchers
  - ◆ Mostly doing research in Clouds, HPC, Big Data
- ▶ Many users need to **deploy Cloud stacks** for their experiments
  - ◆ Typical experiment: compare vanilla vs customized OpenStack
  - ◆ Testbeds should factor this into standard tool or *appliances*
- ▶ Challenging, because:
  - ◆ Cloud stacks are **complex beasts**
  - ◆ **Short release cycles** (6 months) vs staying up-to-date
  - ◆ Need a **low entry barrier** (for tutorials etc.)
  - ◆ Need **support for customization**
  - ◆ Need to **scale** (many-nodes experiments)

# Historical efforts

- ▶ Grid'5000 school, **June 2011**: tutorial about Nimbus and OpenNebula (custom-made scripts)
- ▶ **April 2012**: workshop about *IaaS on Grid'5000*
  - ◆ One solution for OpenStack (custom-made script)
  - ◆ Three solutions for OpenNebula (two using Ruby+Chef, one unspecified)
- ▶ Grid'5000 school, **December 2012**, tutorials:
  - ◆ Nimbus, OpenNebula and Cloudstack (*engines* for an orchestration tool, g5k-campaign)
  - ◆ OpenStack (using PuppetLabs' OpenStack modules + script)
    - ★ Maintained until Grizzly (2013.1)
    - ★ **2014**: Attempts to port it to IceHouse (2014.1) by the technical team, additional problems with Neutron (required 3 NICs)

## Current solution

- ▶ **2015: Users survey:** 10 different ways to deploy OpenStack on Grid'5000 (various versions, various tools)
  - ◆ Most promising user solution made *official*
    - ★ Core: **OpenStack's official Puppet modules**
    - ★ Instantiated on an basic Ubuntu 14.04 image
    - ★ Orchestration using Rake ( $\approx$  Ruby's make)
    - ★ 😊 Liberty and Mitaka supported (complexity in Puppet modules)
    - ★ 😊 Easy to customize (already received users contributions)
    - ★ 😞 Slow to deploy ( $\approx$  30 mins, inc. resources reservation)
- ▶ Related work:
  - ◆ **CloudLab:** One image per node type, Python + bash scripts for setup, Liberty supported, no customization instructions
  - ◆ **Chameleon:** DevStack-based single node deployment, Mitaka supported

# Future challenges

- ▶ **Reduce bootstrap time**
  - ◆ Prepared images? Tradeoff with reproducibility, maintenance effort?
- ▶ **Cost of populating datastores** (VM image repositories, object stores)
  - ◆ Use managed datastores (permanent, testbed-provided)?
    - ★ Tradeoff with reproducibility?
  - ◆ Better strategies for data import?
- ▶ **Extend approach to other frameworks** (e.g. Big Data framework)