

A look at Grid'5000, Clouclab and Chameleon

Lucas Nussbaum
lucas.nussbaum@loria.fr



Grid'5000, Chameleon, CloudLab

▶ **Grid'5000**

- ◆ Established testbed (1000 machines, 500+ users/year since 2005)
- ◆ <https://www.grid5000.fr/>

▶ Two recent projects (2014 ~ 2017; NSF funding for 10 M\$ each)

◆ **Chameleon**

- ★ <https://www.chameleoncloud.org/>

◆ **CloudLab**

- ★ <http://www.cloudlab.us>

Grid'5000, Chameleon, CloudLab

▶ **Grid'5000**

- ◆ Established testbed (1000 machines, 500+ users/year since 2005)
- ◆ <https://www.grid5000.fr/>

▶ Two recent projects (2014 ~ 2017; NSF funding for 10 M\$ each)

◆ **Chameleon**

- ★ <https://www.chameleoncloud.org/>

◆ **CloudLab**

- ★ <http://www.cloudlab.us>

How do they compare: design choices? features?

(To the best of my knowledge, and as of August 2015)

Software stack used as a basis

- ▶ **Grid'5000**: mostly our own (+ OAR, Kadeploy: developed in close collaboration)
- ▶ **Chameleon**: OpenStack
- ▶ **CloudLab**: Emulab

Resources description and verification

- ▶ **Grid'5000**: reference API + g5k-checks (+ human-readable description)
- ▶ **Chameleon**: same as Grid'5000
- ▶ **CloudLab**:
 - ◆ machine-readable description using RSpec 'advertisement' format (less detailed than Grid'5000's, though) + human-readable description in the docs
 - ◆ verification: nothing similar to g5k-checks, but LinkTest¹ can validate network configuration

¹D.S. Anderson et al. "Automatic Online Validation of Network Configuration in the Emulab Network Testbed". In: *ICAC'06*.

Grid'5000 Reference API and g5k-checks²

- ▶ **Describing** resources \leadsto understand results
 - ◆ Covering nodes, network equipment, topology
 - ◆ Machine-parsable format (JSON) \leadsto scripts
 - ◆ Archived (*State of testbed 6 months ago?*)
- ▶ **Verifying** the description
 - ◆ Avoid inaccuracies/errors \leadsto false assumptions, wrong experimental results, retracted publications
 - ◆ Could **happen frequently**: maintenance, broken hardware (e.g. RAM)
 - ◆ Our solution: **g5k-checks**
 - ★ Runs at node boot (or manually by users)
 - ★ Retrieves current description of node
 - ★ Acquires info using OHAI, ethtool, etc.
 - ★ Compares with Reference API

```
"processor": {
  "cache_l2": 8388608,
  "cache_l1": null,
  "model": "Intel Xeon",
  "instruction_set": "",
  "other_description": "",
  "version": "X3440",
  "vendor": "Intel",
  "cache_ll1": null,
  "cache_llid": null,
  "clock_speed": 2530000000.0
},
"uid": "graphene-1",
"type": "node",
"architecture": {
  "platform_type": "x86_64",
  "smt_size": 4,
  "smp_size": 1
},
"main_memory": {
  "ram_size": 17179869184,
  "virtual_size": null
},
"storage_devices": [
  {
    "model": "Hitachi HDS72103",
    "size": 298023223876.953,
    "driver": "ahci",
    "interface": "SATA II",
    "rev": "JPFO",
    "device": "sda"
  }
],
```

²David Margery et al. "Resources Description, Selection, Reservation and Verification on a Large-scale Testbed". In: *TRIDENTCOM. 2014*.

Resources reservation

- ▶ **Grid'5000**: batch scheduler (OAR) with advance reservation
 - ◆ Selection of resources using properties generated from the Reference API
- ▶ **Chameleon**: leases using OpenStack Blazar
- ▶ **CloudLab**: experiments start immediately, default duration of a few hours, can be extended on demand (no advance reservations)

Resources reconfiguration (software level)

All three testbeds have a disk imaging solution:

- ▶ **Grid'5000**: Kadeploy
- ▶ **Chameleon**: OpenStack Ironic
- ▶ **CloudLab**: Emulab's Frisbee

Network reconfiguration and SDN

- ▶ **Grid'5000**: KaVLAN (VLAN reconfiguration on switches, to isolate experiments)
 - ◆ + Distem³ for network emulation
- ▶ **Chameleon**: planned, using OpenFlow
- ▶ **CloudLab**:
 - ◆ Emulab's network emulation features
 - ◆ OpenFlow access on switches⁴
 - ◆ Interconnection to Internet2's AL2S

³<http://distem.gforge.inria.fr>

⁴<http://cloudlab-announce.blogspot.com/2015/06/using-openflow-in-cloudlab.html>

Experiment monitoring

- ▶ **Grid'5000**: Kwapi (power + network) – talk tomorrow
- ▶ **Chameleon**: planned, using OpenStack Ceilometer
- ▶ **CloudLab**: planned⁵ (talk tomorrow?)

⁵<http://docs.cloudlab.us/planned.html>

Long term storage between experiments

- ▶ **Grid'5000**: storage5k (file-based and block-based)
- ▶ **Chameleon**: object store (OpenStack Swift) available soon
- ▶ **CloudLab**: yes⁶, with snapshots (using ZFS) to version data (the snapshots features are not documented yet)

⁶<http://cloudlab-announce.blogspot.fr/2015/04/persistent-dataset.html>

Federation


- ▶ **Grid'5000**: no
 - ◆ In the past, (network) connection to DAS, and to a Japanese testbed
 - ◆ But no real users / use cases
- ▶ **Chameleon**: ? (talk tomorrow)
- ▶ **CloudLab**: Federated with GENI
CloudLab can be used with a GENI account, and vice-versa

Software stacks deployments (OpenStack, Ceph)

- ▶ **Grid'5000**: both OpenStack and Ceph available (but see below)
- ▶ **Chameleon**: ?
- ▶ **CloudLab**: OpenStack available

None of the current solutions work very well. We want:

- ▶ Software stacks **useful to experimenters**
 - ◆ Recent versions (to stay **relevant**)
 - ◆ Easily **customizable** \rightsquigarrow replace components
- ▶ Software stacks **maintainable** in the long run
 - ◆ Despite 6-month release cycles, with a lot of disruptive changes

A photograph of a server room with blue ambient lighting. A person in a light blue shirt and dark pants is standing in the center, reaching into a server rack on the right. The room is filled with rows of server racks, and the floor is a light-colored tile. The lighting is primarily blue, with some white light coming from the ceiling fixtures.

Topics covered:

- ▶ Software stack
- ▶ Resources description and verification
- ▶ Resources reservation
- ▶ Resources reconfiguration
- ▶ Network reconfiguration and SDN
- ▶ Experiment monitoring
- ▶ Long term storage
- ▶ Federation
- ▶ Software stacks deployments

Thanks! Questions?

`lucas.nussbaum@loria.fr`