



University of  
Zurich <sup>UZH</sup>

S3IT: Service and Support for Science IT

---

# Cloud-bursting on SWITCH engines

Sergio Maffioletti

S3IT: Service and Support for Science IT

<http://www.s3it.uzh.ch>

University of Zurich

# Who am I and what I do

Senior **Infrastructure and Application** specialist at University of Zurich.

**S3IT** - Service and Support for ScienceIT:  
<http://www.s3it.uzh.ch>

Worked for more than 10 years in enabling **large scale** scientific applications.

## Toc

- Why UZH/S3IT started to look at SWITCH engines.
- What usecases did we consider.
- How did we do it.
- Important/critical requirements.
- How do we plan to use SWITCH cloud in future.

# Why UZH/S3IT started to look at SWITCH engines

Lack of in-house computational capacity.

Upgrading in-house cloud infrastructure (Hobbes).

<http://www.s3it.uzh.ch/infrastructure/sciencecloud>

Need to satisfy **peak requests** from end-users.

Lot of usecases are computing intensive but can be scaled on a loosely coupled infrastructure.

# What usecases did we consider

Biochemistry

Business

Geography

# What usecases did we consider

## Biochemistry

- Docking Ligands to Proteins with rDock.
- 100 - 200 ligands files to dock.
- Data size negligible.
- On SWITCH cloud: 86 cores x 5 days.

## Business

## Geography

# What usecases did we consider

## Biochemistry

## Business

- Process social network data through Natural Language Processing tools (Stanford NLP tool).
- 20M entries to be processed.
- Data size negligible.
- Data can be processed independently.
- On SWITCH cloud: 100 cores x 5 days.

## Geography

# What usecases did we consider

Biochemistry

Business

Geography

- Process GPS data.
- 1 RData file to be split in grid regions.
- Each execution analyses 1 grid area.
- Grids processed independently.
- 50 grids to be processed.
- Data size negligible.
- On SWITCH cloud: 50 cores x 2 days.



## How we did it: Bursting from Hobbes

Migrate or re-create reference images

Re-create same resource configuration: VMs and data volumes

Adapt workflows

## How we did it: Bursting from Hobbes

Migrate or re-create reference images

- [glance tools](#) to copy existing images.
- [Ansible playbooks](#) to re-create images from vanilla Linux.

Re-create same resource configuration: VMs and data volumes

Adapt workflows

## How we did it: Bursting from Hobbes

Migrate or re-create reference images

Re-create same resource configuration: VMs and data volumes

- [elasticcluster](#) for creating dedicated SLURM clusters.
- <http://gc3-uzh-ch.github.io/elasticcluster/>
- Ansible playbooks underneath.

Adapt workflows

## How we did it: Bursting from Hobbes

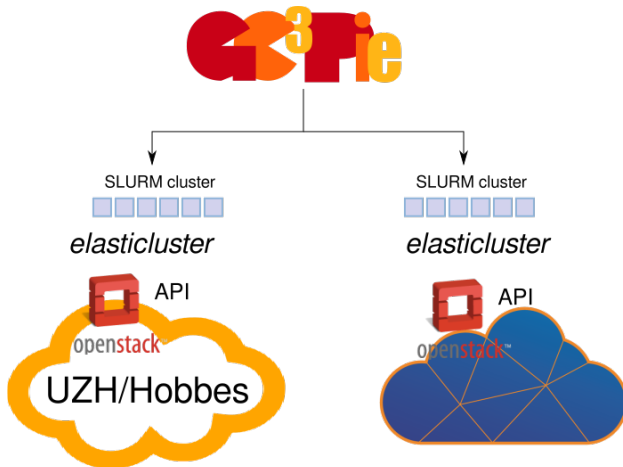
Migrate or re-create reference images

Re-create same resource configuration: VMs and data volumes

Adapt workflows

- GC3Pie: <https://code.google.com/p/gc3pie/>
- Can distribute workload on multiple resources.
- Simple change in configuration file on client side.

# How we did it: Bursting from Hobbes



## Important/critical requirements

IaaS **API** compatibility (OpenStack and/or EC2).

Using existing tools seamlessly.

Use same/similar procedures for creating images, volumes and networks.

Fast allocation of **bulk of resources**.

## Future perspectives

**Server hosting:** need to have clear SLAs and guaranteed uptime.

Continue in **cloud bursting** - data transfer is the next question.

**Educational** purposes - training classes and/or MOOCs (e.g. edx).