University of Zurich[UZH]

S^{3}IT: Services and Support for Science IT

What is ScienceCloud ?

Antonio Messina <antonio.messina@uzh.ch>
Cloud Architect
Sergio Maffioletti <sergio.maffioletti@uzh.ch>
User support
University of Zurich, S^{3}IT

23 November 2015

Why did we build ScienceCloud ?

Store, access and process research data.

Facilitate UZH research infrastructure investment planning.

S3IT research infrastructure strategy:

- High Performance Computing
- Cluster Computing
- Server Computing
- Research Data Storage

What criteria have been driving this ?

Flexibility for end-users.

Self-provisioning / elasticity of VMs, Storage and Network (ultimately of Services).

Scalability and extensibility of the underlying infrastructure.

Reliability and availability of the services.

Why would you care?

Researchers' FAQ:

- How can I run this data analysis on 1000 cores since on my laptop is too slow? (btw, I need to submit for publication by end of this month)
- Where can I put this 100TB of data that I need to analyze? (did I tell you I have a deadline end of this month?)
- How can I automate all of this? Can you do it for me?
- Do I need to adapt my application to run on your system? Can you do it for me?

The infrastructure adapts to the use case

ScienceCloud is not an isolated service.

S^{3}IT provides solutions for your data analysis usecase:

- Usecase analysis
- Solution enginneering and implementaion
- tools to run large scale data alaysis and to automate the infrastructure provisioning:
  - GC3Pie
  - Elasticluster
- Development to implement large-scale data analysis solutions
- and the infrastructure where to run it.

Why a cloud infrastructure then ?

Self-provisioning and Elasticity of resources

Customization and control of the environment

Network API

Why a cloud infrastructure then ?

Self-provisioning and Elasticity of resources

- End-users can allocate and release resources when needed.

Customization and control of the environment

Network API

Why a cloud infrastructure then ?

Self-provisioning and Elasticity of resources

Customization and control of the environment

- End-users can tailor the research infrastructure to their specific needs.

Network API

Why a cloud infrastructure then ?

Self-provisioning and Elasticity of resources

Customization and control of the environment

Network API

- To programmatically create and control an own research infrastructure.
- GC3Pie: http://gc3pie.readthedocs.org/en/latest/
- Elasticluster: http://gc3-uzh-ch.github.io/elasticluster/
- custom *glue* scripts

Design requirements

An University wide infrastructure like this needs to be

- flexible
- horizontally scalable
- reliable
- fast

Speed

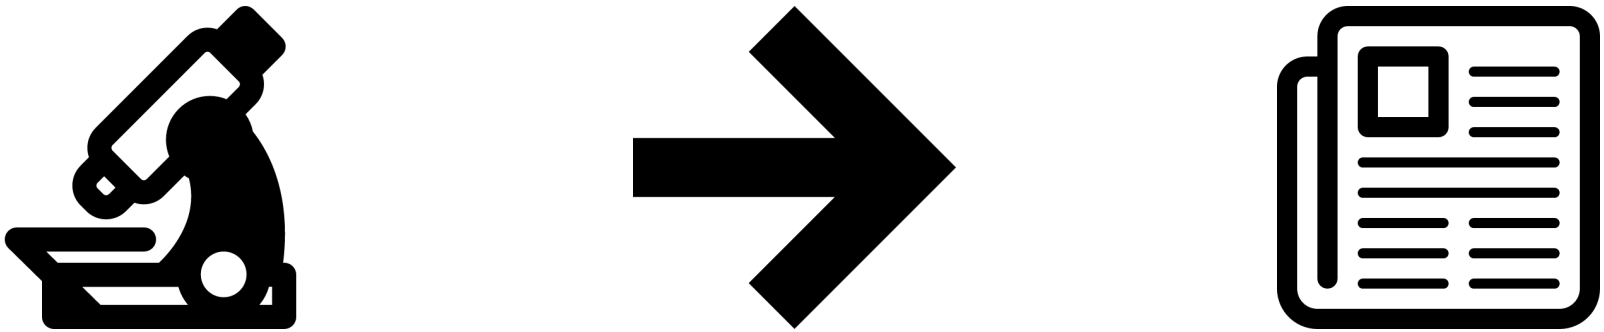

speed is ***time to solution***

- Fast startup of VMs
- Support in design/implement of the infrastructure
- Fast cpu/network/storage

Training/support

- help@s3it.uzh.ch for any request/suggestion
- Online documentation on the S^{3}IT wiki (please provide feedback!)
- One day per month allocated for ScienceCloud training (starting from November 2015.)
- Ad-hoc training and site visits on demand

Hardware

Now:

- 1536 cores with HyperThreading
  - 6 blade chassis
  - 96 dual-socket E5-2640 v3 16 cores (Haswell), 128GB RAM each
- 12TB RAM total
- 1PB **usable** storage
- 36 storage nodes

End of 2015:

- 3648 cores (7296 threads)
- 19TB RAM total
- 1.7 PB