



Shaping HPC cloud services
experiences and lessons

tom.visser@sara.nl

[linkedin.com/in/visservisser](https://www.linkedin.com/in/visservisser)

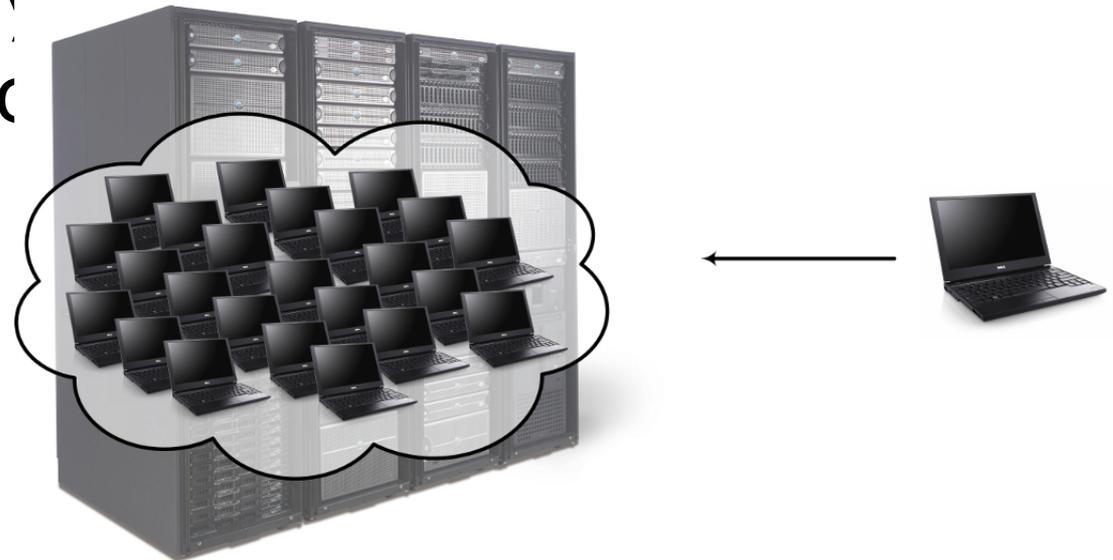
BiG Grid

the dutch e-science grid

How it started

Cloud computing is not about new technology, it is about new uses of technology

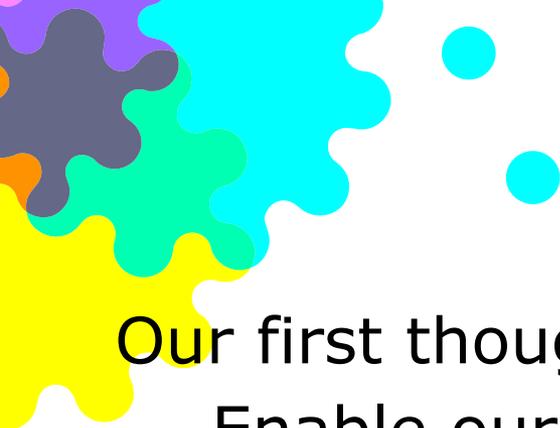
Self Service Dynamically Scalable
Computing Facility



BiG Grid

the dutch e-science grid





Or did it?

Our first thoughts:

- Enable our users to do new things
- **Change** the way we do things i.e. operate our datacenters and offer support
 - rent hardware
 - consolidate & integrate all infra.
 - specialize in specific services
 - change support model
- Virtualize all hardware
 - we did that with our grid-server park



BiG Grid

the dutch e-science grid



The pilot phase 2009

When wanting to build something new, challenge yourself to directly involve the users!

Open concrete invitation i.e. call for experimentation using our crowd:

2 months only, **intensive personal contact** with everyone

Lessons:

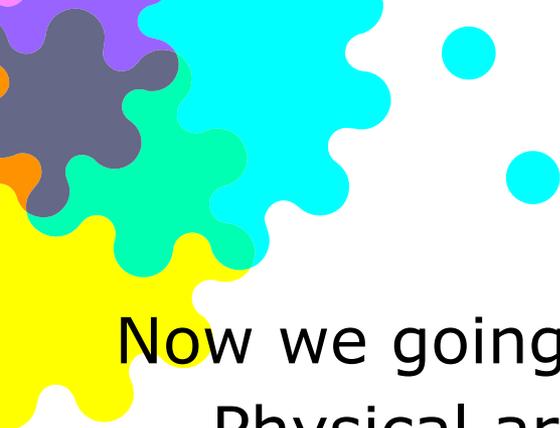
- Flexibility and shorter time to solution
- Massive discussions about security
- You **can** do everything yourself but you also

have to

Supporting the supporter: i.e. enabling scientific

programmers





Pre-production 2010-2011

Now we going for the real thing.

- Physical architecture, Usability, Security
- Offering users their own meta-support portal (support the supporter)

Longer phase; no formal application procedure

Anything *but* hosting websites

- Dynamic R clusters
- Taverna language processing and on demand
- Experimenting with occi and sniaa standards; built cdmfs
- Built our own user control portal

It just worked.



BIG Grid

the dutch e.science grid

Production

- 400TB NFS mountable; 19 x 32 core nodes, 256 GB mem.
- Using, patching and modifying Opennebula Experience
- Formal procedure, still 2 week pilot accounts
- Supporting the university IT-departments
- What is our roadmap, can we still experiment?
- Fast growth
- Keep on publicly communicating
- Keep users involved
- Duplication of core services?
- Slowly moving up from not only IAAS _ SIAAS _

Thoughts - Dilemmas - outlook

What will a federation of cloud offer scientific users?

- are we building another grid?
- wall clock time, single instances
- homogeneous

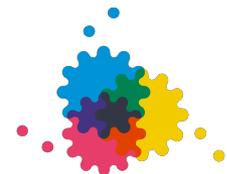
Do we need a vision; or just an approach

We are being capable of supporting more and different scientific usage patterns with less people

Consolidation needs scheduling taxonomy

for example: auto scaling, never scaling, run until

Strong ecosystem! ; compatible and diverse



BiG Grid

the dutch e-science grid



Great ecosystem e.g:

<http://www.datawrangling.com/mpi-cluster-with-python-and-amazon-ec2-part-2-of-3>

<http://www.drewconway.com/zia/?p=2701> EC2 AMI for scientific computing in Python and R

Genomics: <http://cloudbiolinux.org/>

Galaxy workflows using <http://usecloudman.org/> and Galaxy cloudman



BiG Grid

the dutch e-science grid

One final lesson



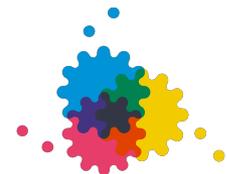
Be grateful!

Thanks to:

All of our team at SARA coworkers of the last 2,5 years, and BiG Grid, all of our enthusiastic beta users. You all know who you are. All developers. EGI for having this opportunity to talk. You!

Special mention:

Floris Sluiter, the inspiring project lead for the last 2 years.



BiG Grid

the dutch e-science grid