BiG Grid HPC Cloud^{Beta}

Floris Sluiter
SARA Computing and Networking services
Amsterdam
www.cloud.sara.nl

BiG Grid
the dutch e·science grid

About BiG Grid and Sara

The BiG Grid project is a collaboration between NCF, Nikhef and NBIC, and enables access to grid infrastructures for scientific research in the Netherlands.

SARA is a **national** High Performance Computing and e-Science **Support Center**, in Amsterdam and the primary operational partner of BiG Grid

sara

BiG Grid
the dutch e-science grid

"Our" definition of Cloud

Cloud Computing:

Self Service Dynamically Scalable Computing Facilities

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

Vision: Clone my laptop!

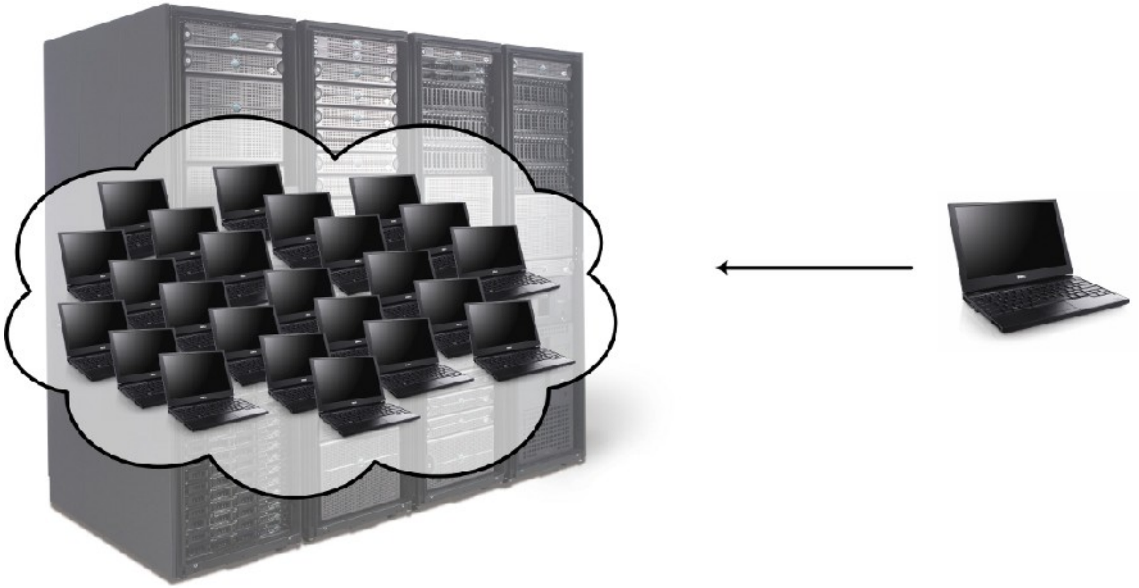

BiG Grid
the dutch e-science grid

Virtual Private HPC Cluster

Fully configurable HPC Cluster

Freedom of choice

Secure environment

Self service GUI

Virtual Private HPC Cluster 1

Fully configurable HPC Cluster for every user

- Fast CPU
- Large Memory (64GB/8 cores)
- High Bandwidth network (40Gbit/s Infiniband)
- Large and fast storage

Virtual Private HPC Cluster 2

Freedom of choice

- Users will be root inside their own cluster
- Built cluster from scratch
- Free choice of OS, etc
- And/Or use existing VMs: Examples, Templates, Clones of Laptop, Downloaded VMs, etc
- Public IP possible (subject to security scan)

Virtual Private HPC Cluster 3

Security

- Separated network, each VPHC in own Vlan
- Firewalled (self service)
- Network scanning
- etc...

Virtual Private HPC Cluster 4

Self service GUI

- Web based
- Access with modern browser
- Developed in house (open sourced)

Development Roadmap

- 2009, Q3 Q4: Pilot Phase (finished)
    - Small testbed, 50 cores, 5 usergroups
- 2010, Q2, Q3: Pre-production Phase (**almost** finished)
    - Medium sized testbed, 128 cores, 100 Tbyte storage
- 2010, Q4,Q++: Production Phase
    - >=1024 cores *planned,* configuration pending

Follow Standards: OCCI & CDMI

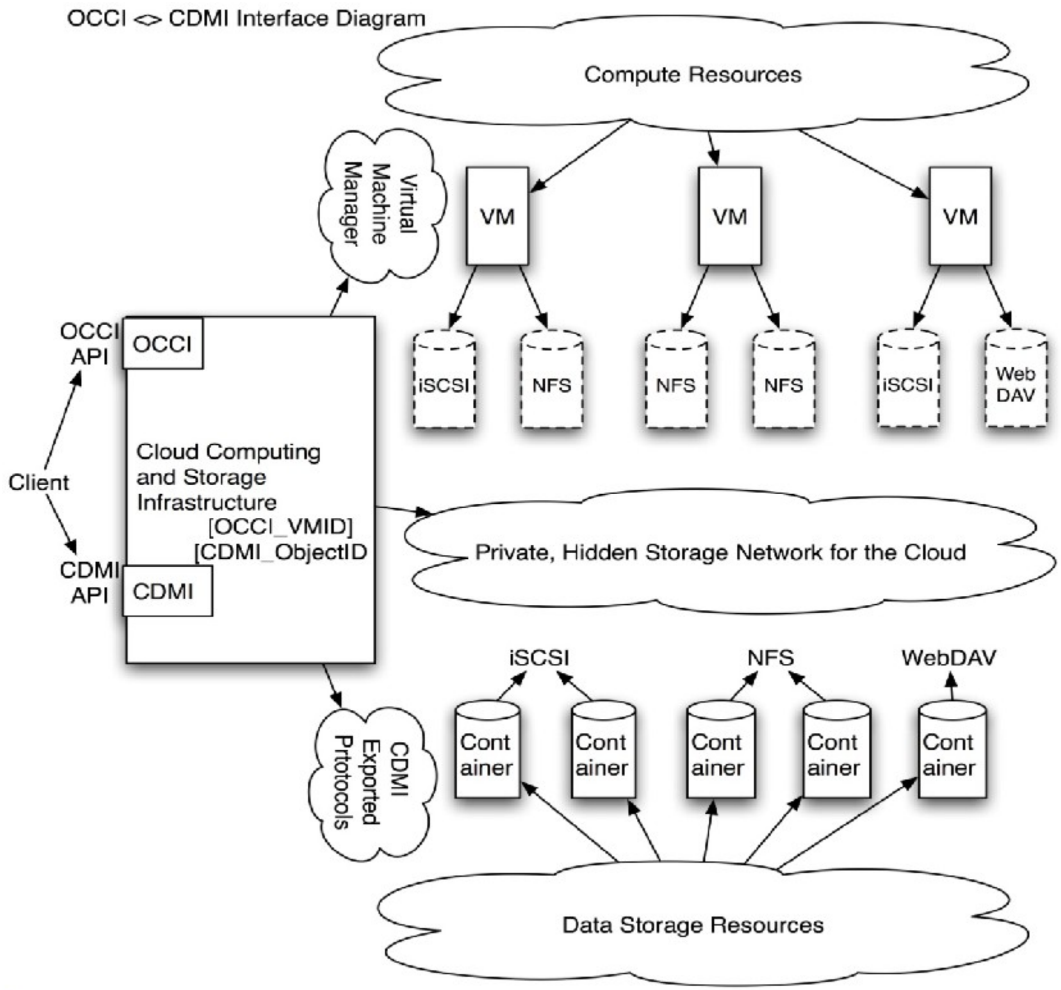

BiG Grid
the dutch e-science grid

A bit of Hard Labour

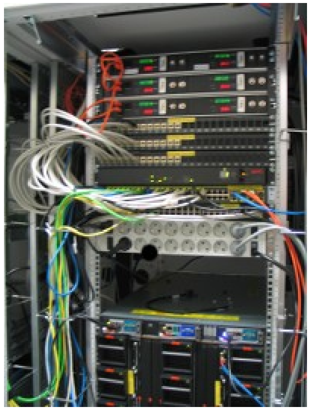

BiG Grid
the dutch e-science grid

Self Service GUI

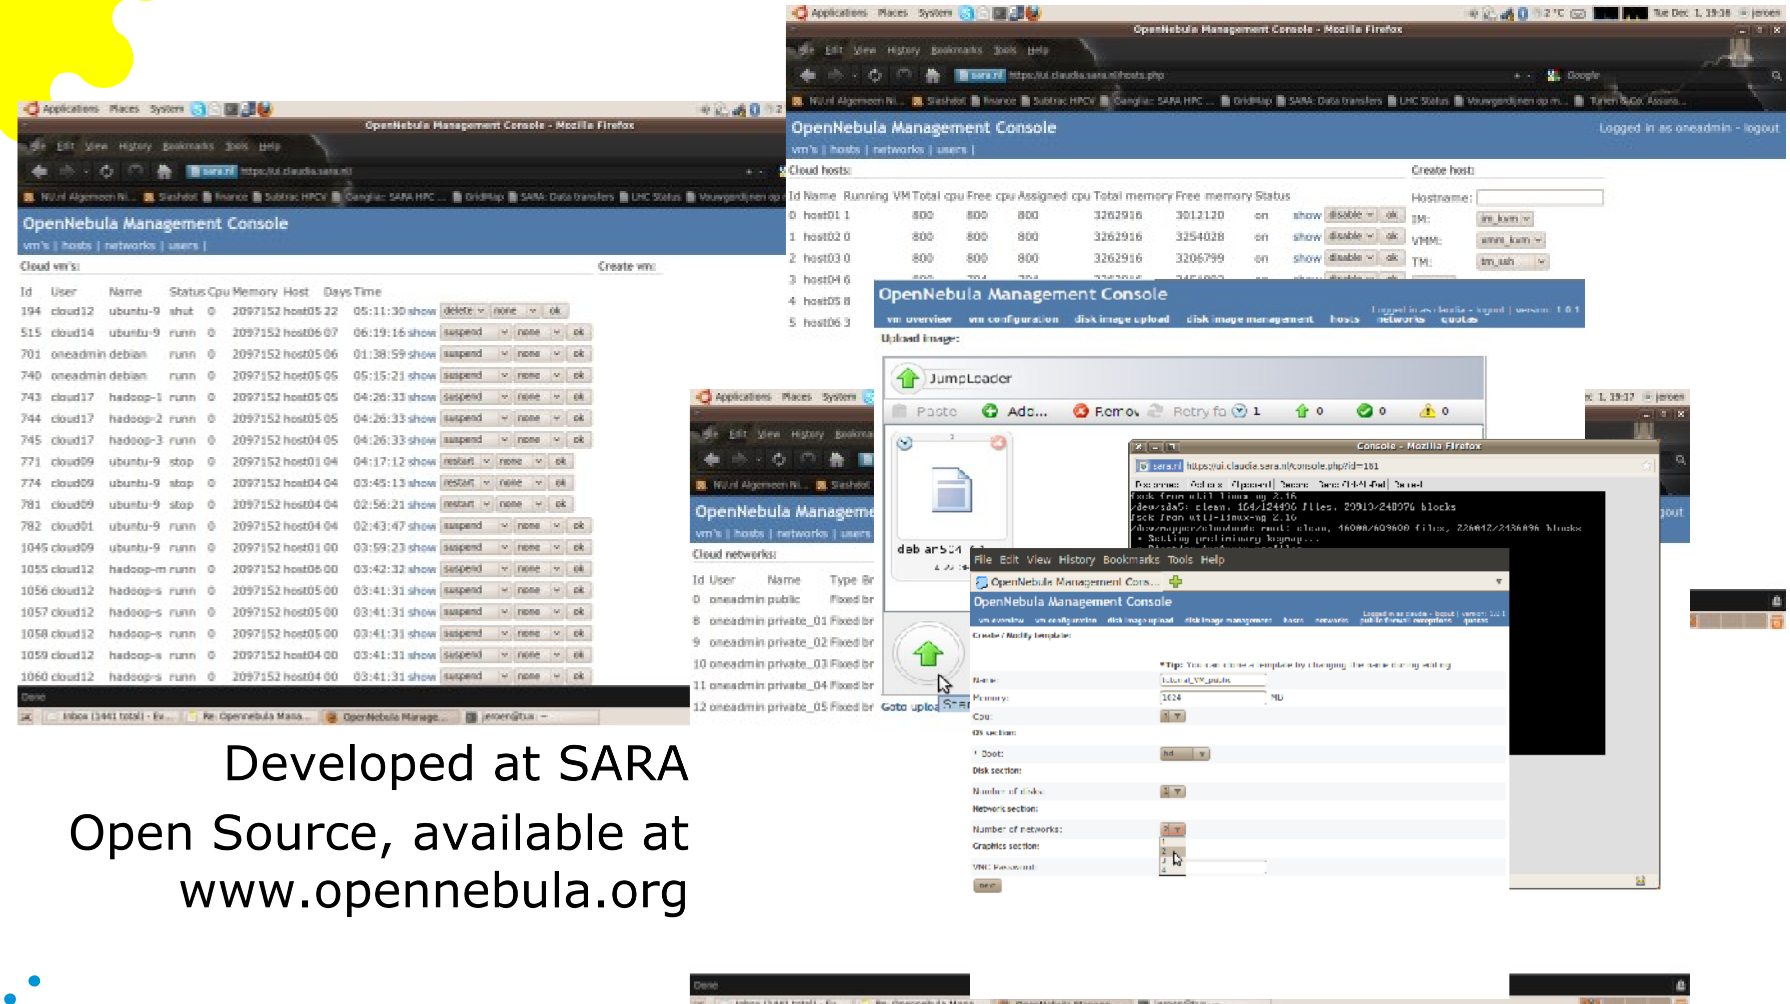

Developed at SARA

Open Source, available at www.opennebula.org

BiG Grid

the dutch e-science grid

Pre-production Phase
From POC to Pr.E...

- Physical Architecture
- Usability
- Security

Pre-production Phase 1
From POC to Pr.E…

- **Physical Architecture**
  - Performance tuning: optimize hard- & software
  - Scheduling
  - HPC Cloud needs High I/O capabilities

Pre-production Phase 2
From POC to Pr.E...

- **Usability**
  - Interfaces
  - Templates
  - Documentation & Education
  - Involve **users** in pre-production (!)

Pre-production Phase 3
From POC to Pr.E...

- **Security**
  - Protect user against self, fellow users, the world and vice versa!
  - Enable user to share private data and templates
  - Self Service Interface
    - User specifies "normal network traffic", ACLs & Firewall rules
  - Monitoring, Monitoring, Monitoring!
    - No control over contents of VM
    - monitor its ports, network and communication patterns

Physical architecture
in this phase

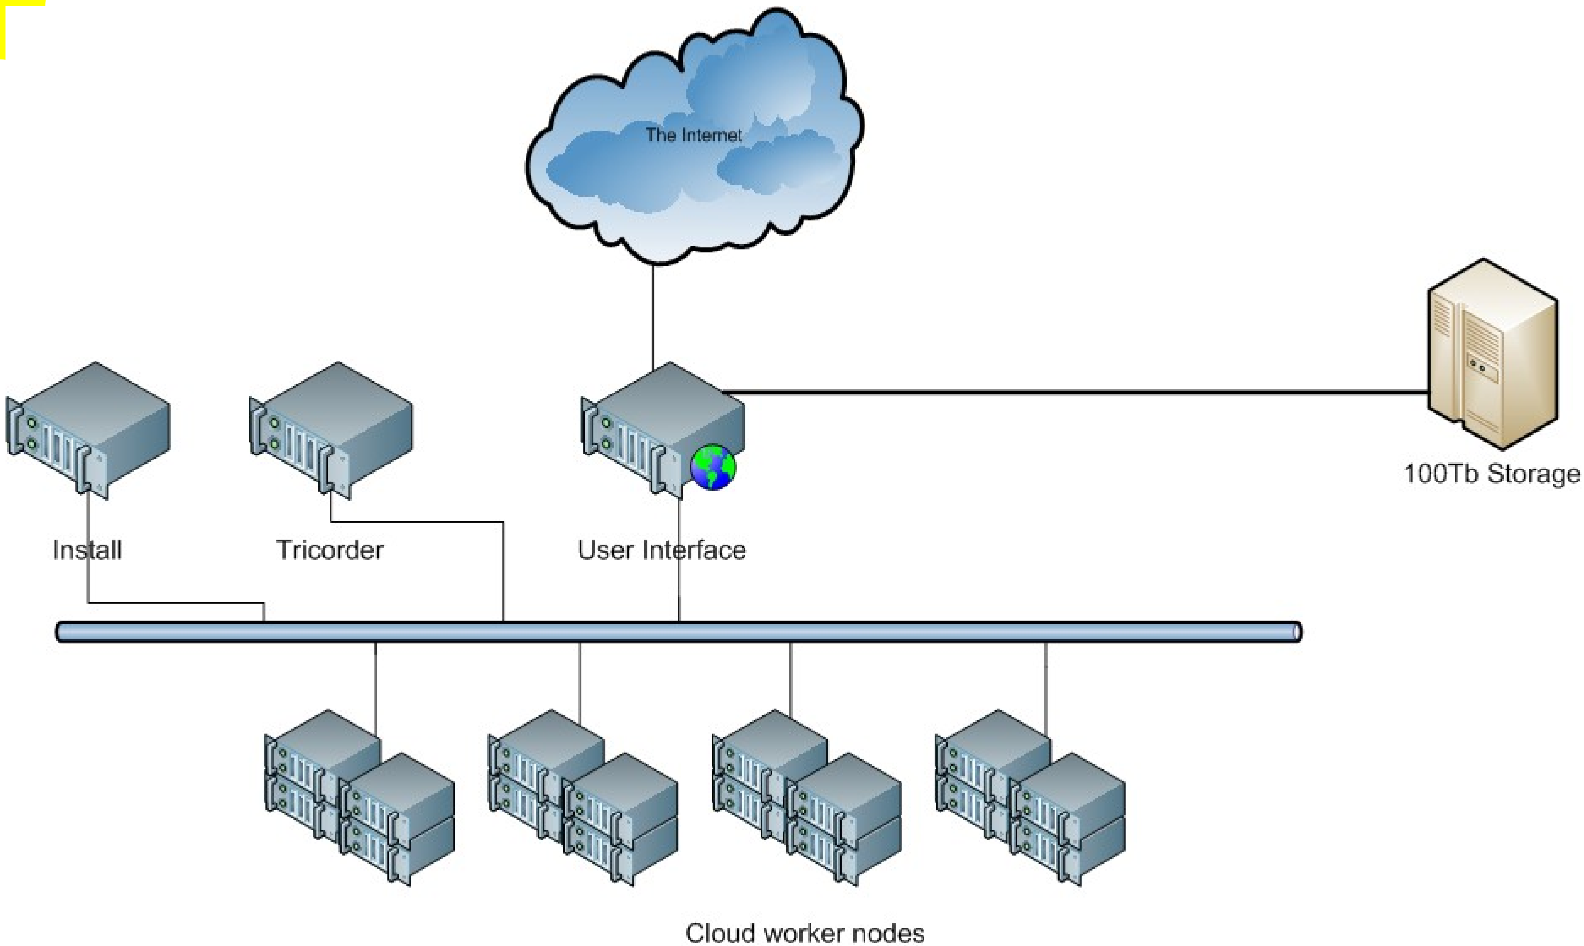

BiG Grid
the dutch e-science grid

Virtual architecture

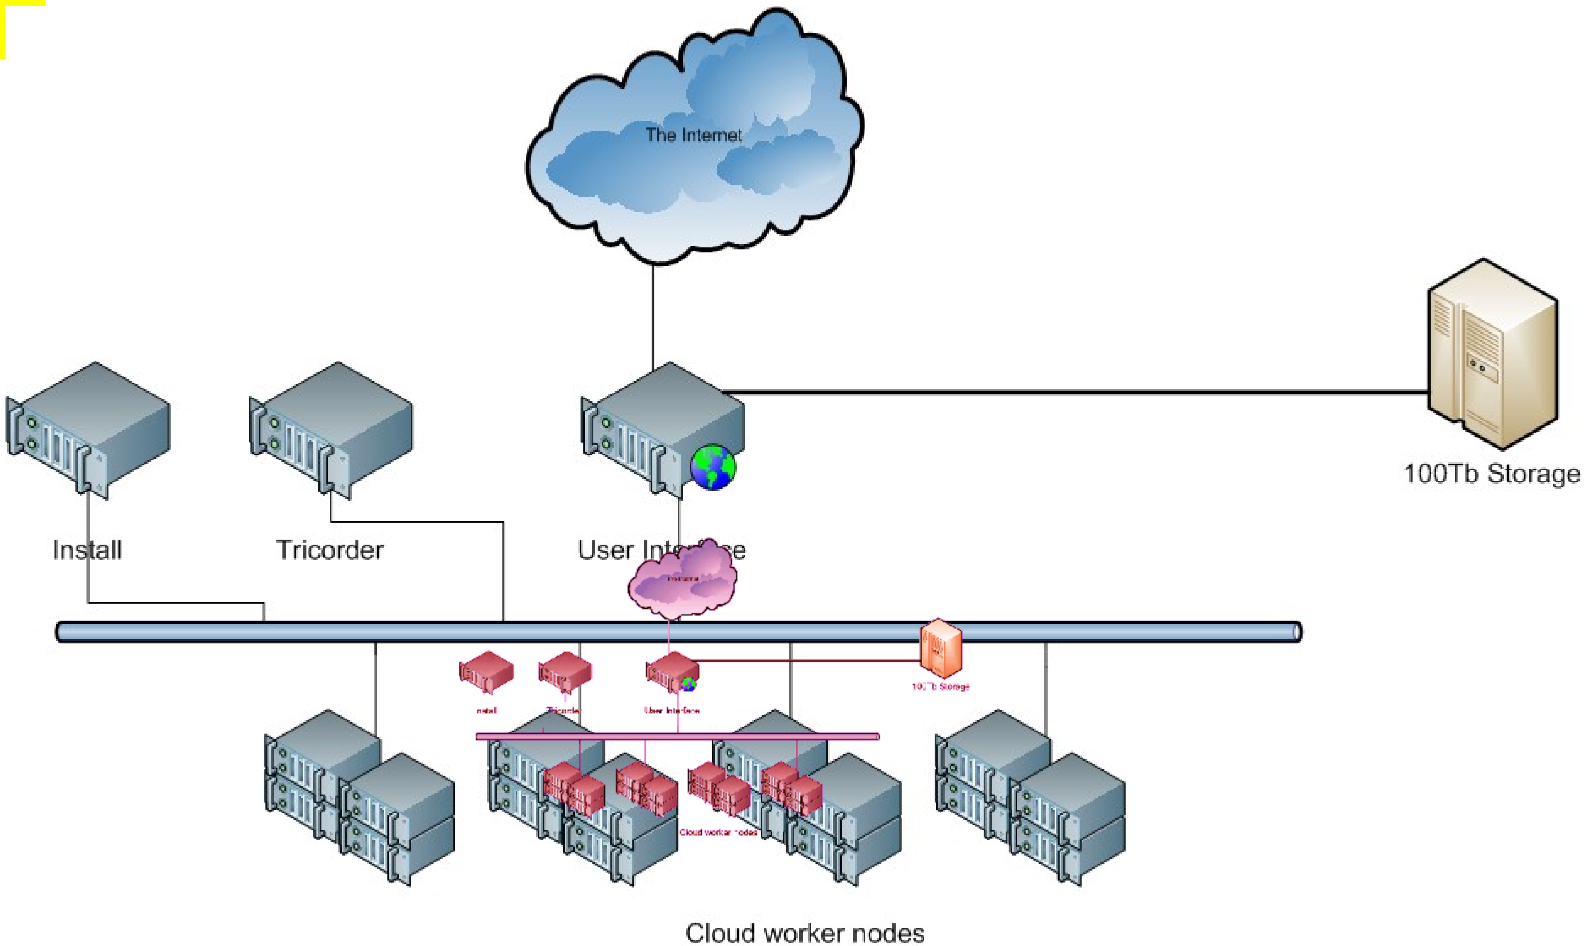

BiG Grid
the dutch e-science grid

Virtual architecture cont...

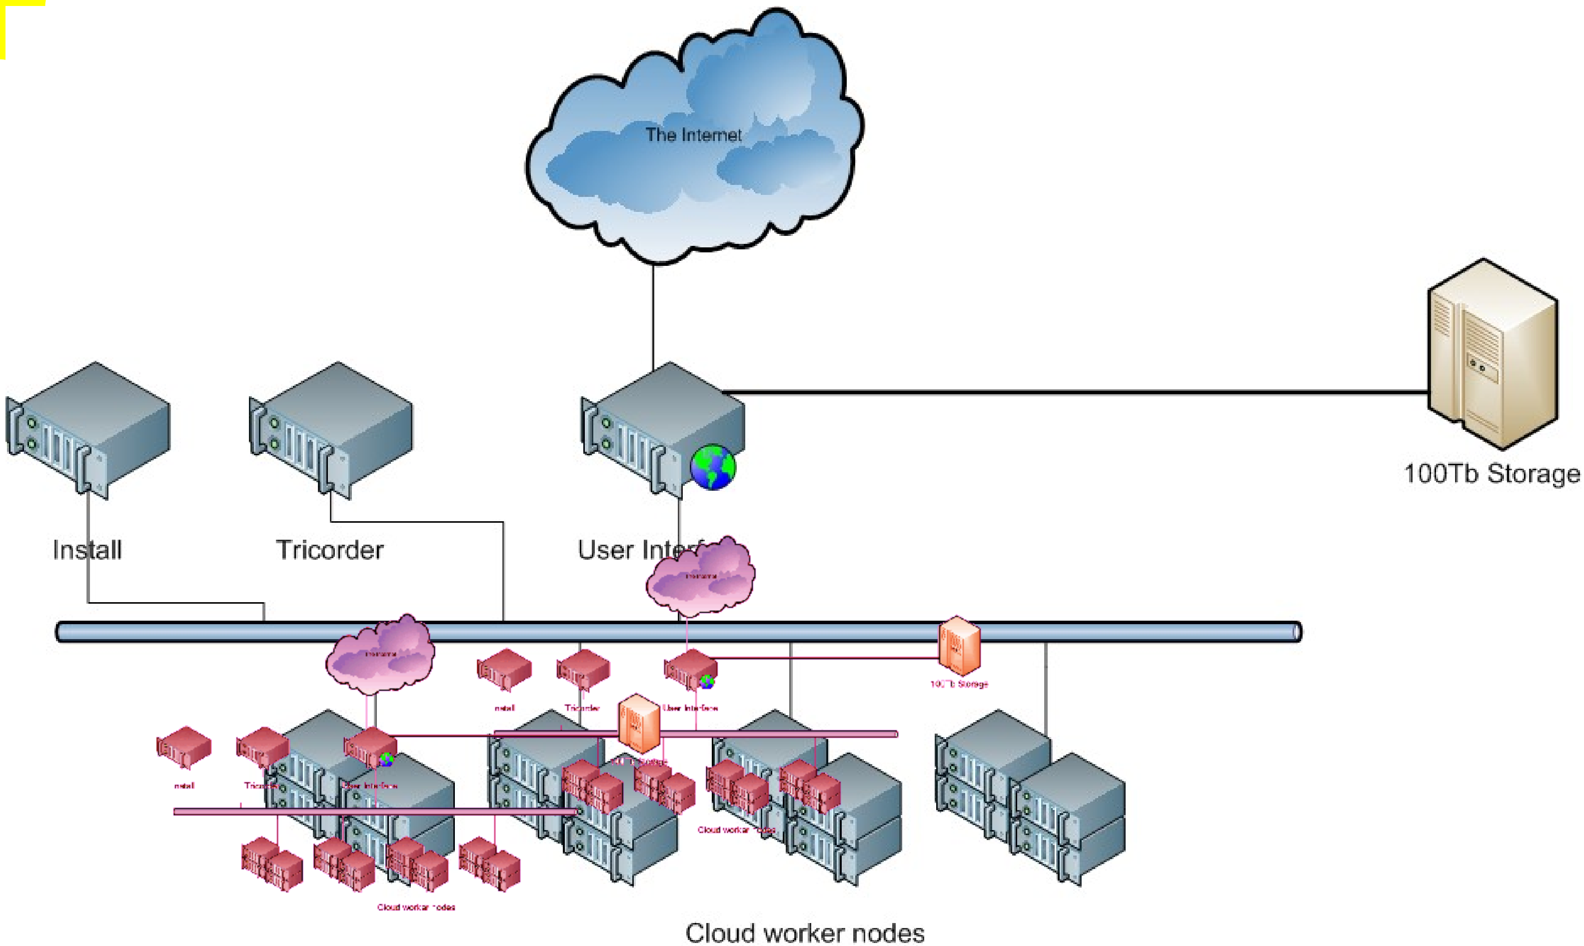

BiG Grid
the dutch e-science grid

Virtual architecture cont...

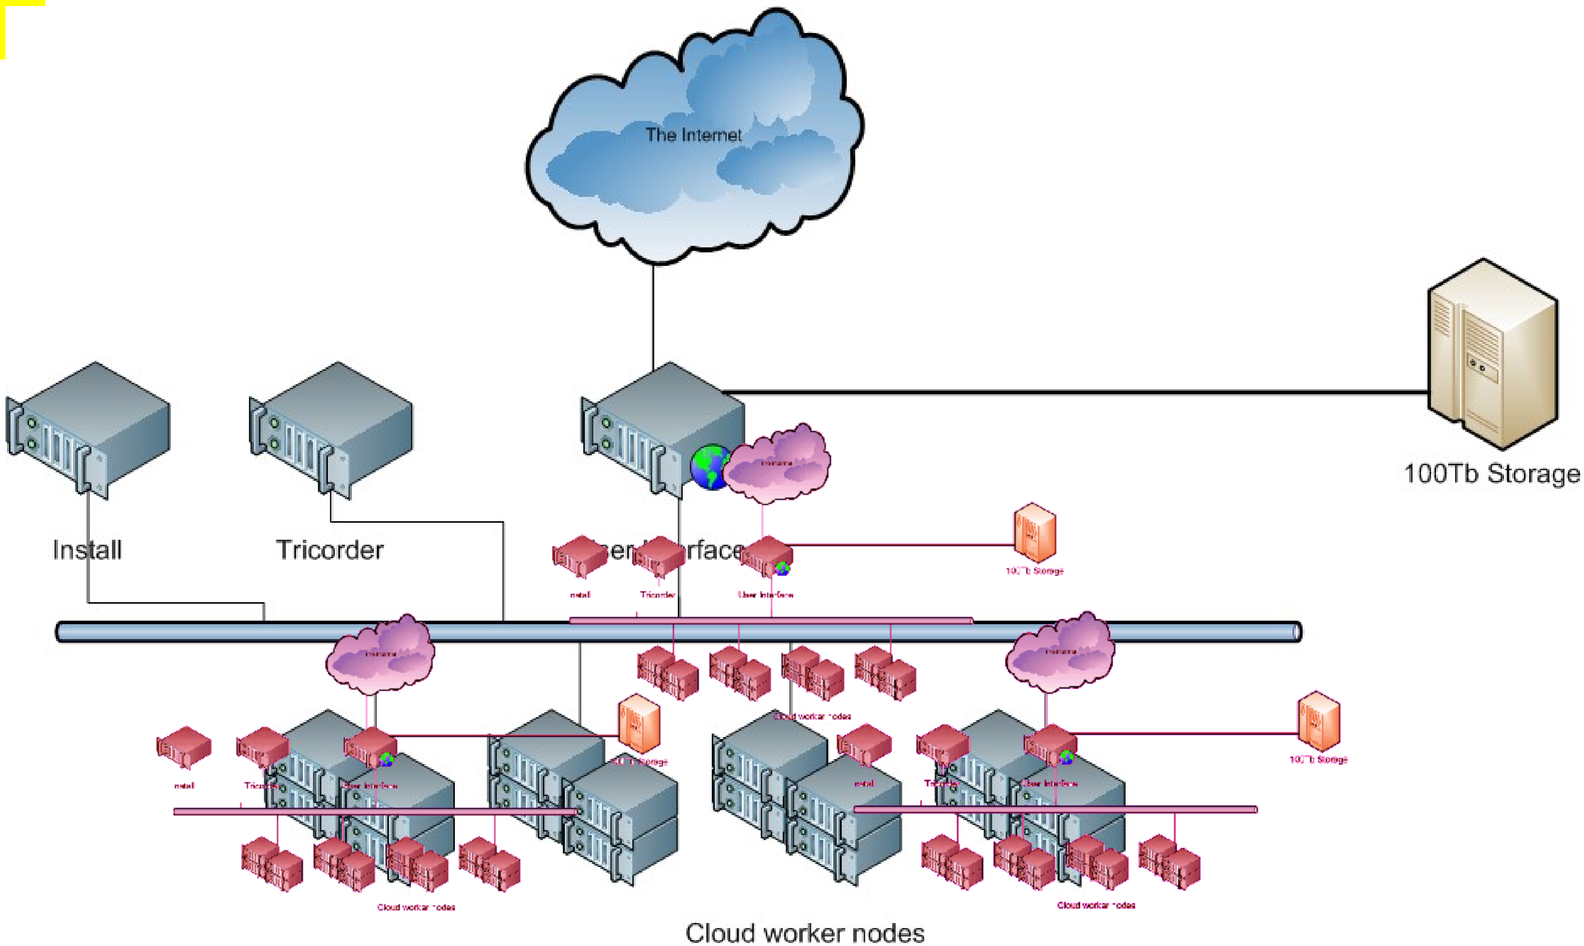

Virtual architecture cont...

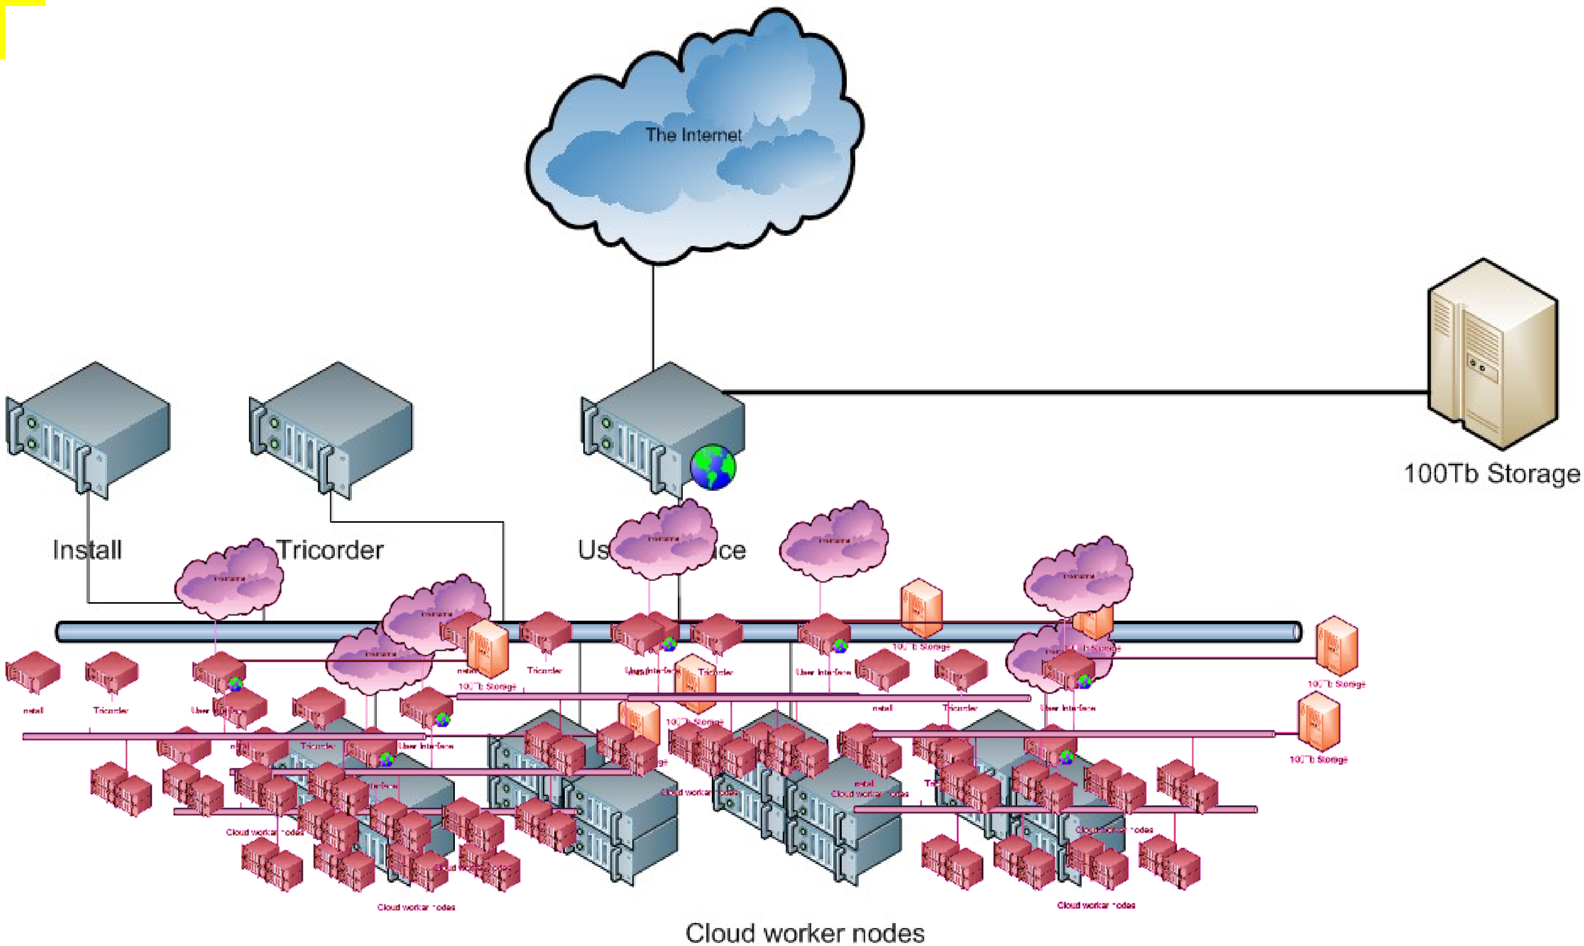

BiG Grid
the dutch e-science grid

Being a pioneer is fun …

Expert Administrators/developers
to develop the infrastructure
(and users do not notice the complexity)!!!

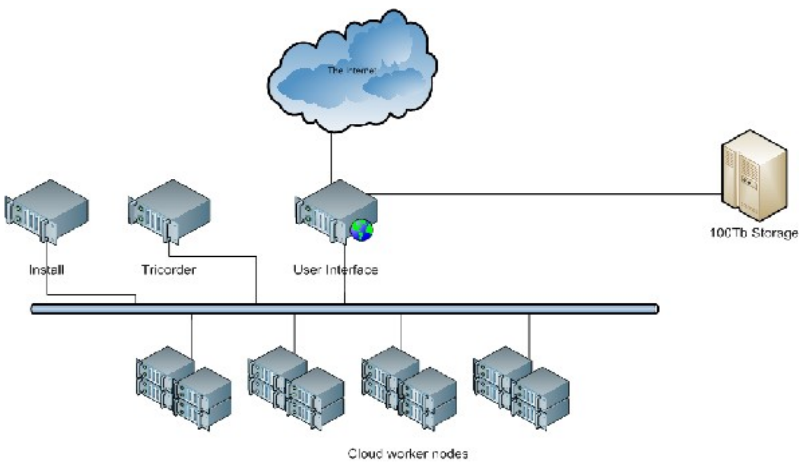

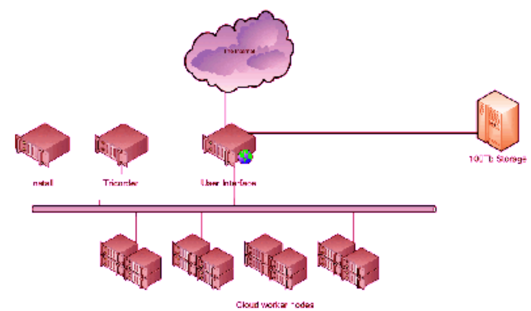

Application types

Type	Examples	Requirements
Compute Intensive	Monte Carlo simulations and parameter optimizations, etc	CPU Cycles
Data intensive	Signal/Image processing in Astronomy, Remote Sensing, Medical Imaging, DNA matching, Pattern matching, etc	I/O to data (SAN File Servers)
Communication intensive	Particle Physics, MPI, etc	Fast interconnect network
Memory intensive	DNA assembly, etc	Large (Shared) RAM
Continuous services	Databases, webservers, webservices	Dynamically scalable

User participation
12 involved in Beta testing

nr.	Title	Core Hours	Storage	Objective	Group/instiute
1	Cloud computing for sequence assembly	14 samples * 2 vms * 2-4 cores * 2 days = 5000	10-100GB / VM	Run a set of prepared vm's for different and specific sequence assembly tasks	Bacterial Genomics, CMBI Nijmegen
2	Cloud computing for a multi-method perspective study of construction of (cyber)space and place	2000 (+)	75-100GB	Analyse 20 million Flickr Geocoded data points	Uva, GPIO institute
3	Urban Flood Simulation	1500	1 GB	asses cloud technology potential and efficiency on ported Urban Flood simulation modules	UvA, Computational Science
4	A user friendly cloud-based inverse modelling environment	testing	1GB / VM	running in the cloud supporting modelling, testing and large scale running of model.	Computational Geo-ecology, UvA
5	Real life HPC cloud computing experiences for MicroArray analyses	8000	150GB	Test, development and acquire real life experiences using vm's for microarray analysis	Microarray Department, Integrative BioInformatics Unit, UvA
6	Customized pipelines for the processing of MRI brain data	?	up to 1TB of data -> transferred out quickly.	Configure a customized virtual infrastructure for MRI image processing pipelines	Biomedical Imaging Group, Rotterdam, Erasmus MC
7	Cloud computing for historical map collections: access and georeferencing	?	7VM's of 500 GB = 3.5 TB	Set up distributed, decentralized autonomous georeferencing data delivery system.	Department of Geography, UvA
8	Parallellization of MT3DMS for modeling contaminant transport at large scale	64 cores, schaling experiments / * 80 hours = 5000 hours	1 TB	Goal, investigate massive parallell scaling for code speed-up	Deltares
9	An imputation pipeline on Grid Gain		20TB	pipelines and, in particular, heavy imputation pipelines on a new HPC cloud	Groningen Bioinformatics Center, university of groningen
10	Regional Atmospheric Soaring Prediction	320	20GB	Demonstrate how cloud computing eliminates porting problems.	Computational Geo-ecology, UvA
11	Extraction of Social Signals from video	160	630GB	Video Feature extraction	Pattern Recognition Laboratory, TU Delft
12	sequencing data from mouse tumors	?	150-300GB	Run analysis pipeline to create mouse model for genome analysis	Chris Klijn, NKI

User participation
12 involved in Beta testing

Field	# projects
Bioinformatics	5
Ecology	2
Geography	3
Computer science	2

User Experience

(slides from Han Rauwerda, transcriptomics UVA)

Microarray analysis: *Calculation of F-values in a 36 * 135 k transcriptomics study using of 5000 permutations on 16 cores.*

worked out of the box (including the standard cluster logic)

no indication of large overhead

Ageing study - *conditional correlation*

dr. Martijs Jonker (MAD/IBU), prof. van Steeg (RIVM), prof. dr. v.d. Horst en prof.dr. Hoeymakers (EMC)
- 6 timepoints, 4 tissues, 3 replicates and 35 k measurements + pathological data
- Question: find per-gene correlation with pathological data (staining)
- Spearman Correlation conditional on chronological age (not normal)
- p-values through 10k permutations (**4000 core hours** / tissue)

Co-expression network analysis
- 6k * 6k correlation matrix (conditional on chronological age)
- calculation of this matrix parallellized. (**5.000 core hours** / tissue)

Development during testing period (real life!)

Conclusions

- Many ideas were tried (clusters with 32 - 64 cores)
- Cloud cluster: like a real cluster
- Virtually no hick-ups of the system, no waiting times
- User: it is a very convenient system

Usage statistics in beta phase

- Users liked it:
  - ~90.000 core-hours used in 10 weeks (~175.000 available)
  - 50% occupation *during beta testing*
  - Some pioneers paved the way for the rest ("Google" launch approach)
  - Evaluation meeting with users, outcome was very positive

Advantages of HPC Cloud

- Only small overhead from virtualization (5%)
- easy/no porting of applications
- Applications with different requirements can co-exist on the same physical host
- Long running services (for example databases)
- Tailored Computing
- Service Cost shifts from manpower to infrastructure
- Usage cost *in HPC* stays Pay per Use
- Time to solution shortens for many users

Acknowledgements

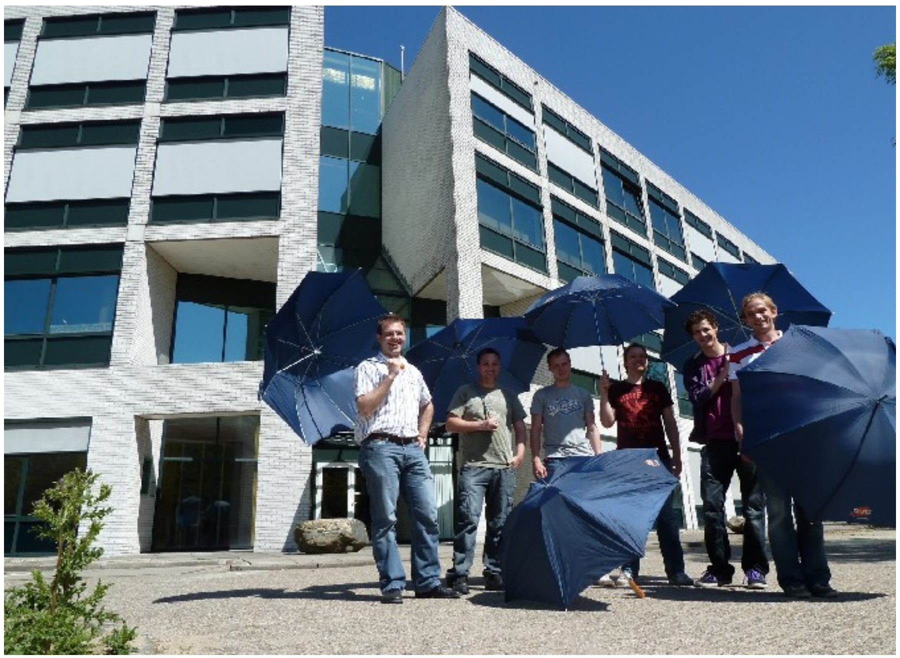

Our Sponsor: NL-BiGGrid

Our Brave & Entrepeneurial **Beta Users**

And the HPC Cloud team:

Tom Visser, Neil Mooney, Jeroen Nijhof, Jhon Masschelein, Dennis Blommesteijn, et. al.

http://www.cloud.sara.nl

photo: http://cloudappreciationsociety.org/

Thank you!

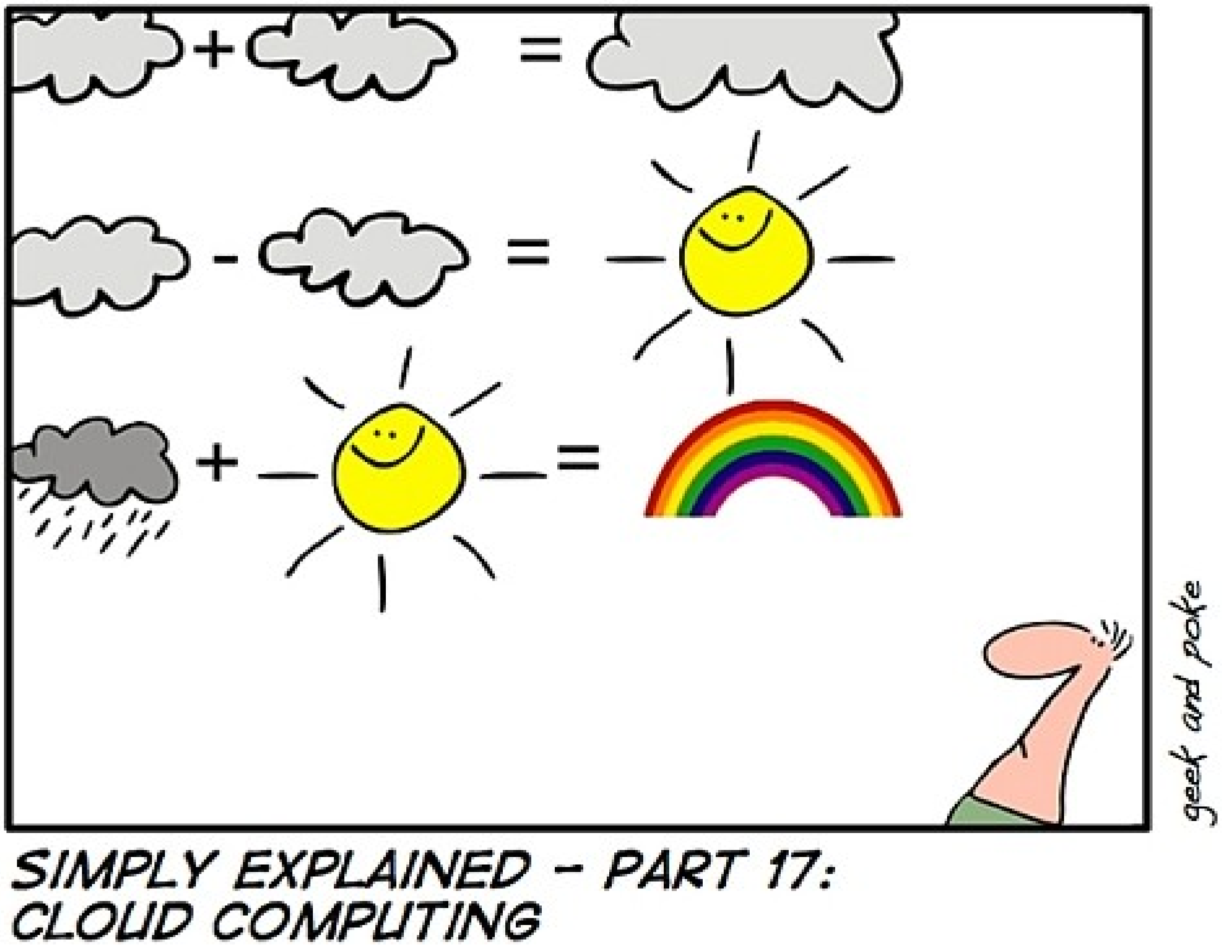

Questions?

BiG Grid
the dutch e-science grid

What else is Cooking?

Extra features:

- AAA
    - Sharing resources
    - Accounting also on I/O & infra
    - Ldap / x509
- Finegrained firewall
- Scheduling also on memory and i/o bandwidth
- **Selve Service Storage**
    - **CDMIFUSE (prototype = working)**
- **Self service networking**
    - Please supply use cases!
- More experiments!

BiG Grid
the dutch e-science grid