HPC Cloud
Calligo system

“Focus on your research”

Floris Sluiter
Project leader
SARA

SARA Project involvements

40 JAAR 1971 2011 sara

BiG Grid
the dutch e-science grid

vl·e

DEISA
DISTRIBUTED EUROPEAN INFRASTRUCTURE FOR SUPERCOMPUTING APPLICATIONS

SURF NET

PRACE

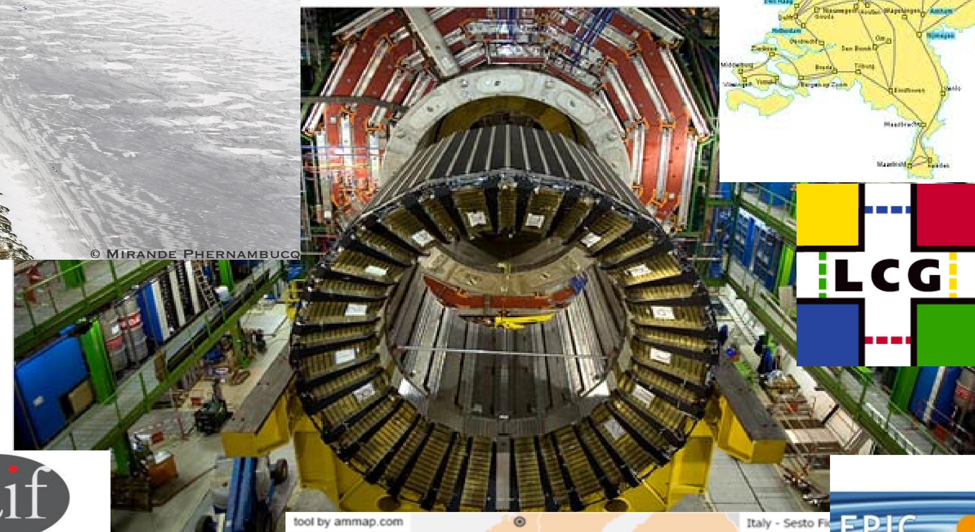

LCG

egi European Grid Infrastructure
Towards a sustainable grid infrastructure

AMSTERDAM LIGHTHOUSE

CINEGRID AMSTERDAM

glif
Global Lambda Integrated Facility

EPIC

HPDMnet
High Performance Digital Media Network

BBMRI
Biobanking and Biomolecular Resources Research Infrastructure

HPC-Europa2
Pan-European Research Infrastructure on High Performance Computing

GÉANT

OptIPuter

nbic

What is High Performance Cloud Computing?

- Cloud Computing: Self Service Dynamically Scalable Computing Facilities
  - Cloud computing is not about new technology, it is about new uses of technology
- “HPC Cloud” is a “Cloud built with high end HPC hardware”
- Through virtualization:
  - Better seperation between users, allowing more administrative rights
  - Very fast to set up custom environments, saving application porting time
  - Users can fully define and control a private HPC cluster

Why an HPC *Cloud*?

- Christophe Blanchet, IDB - Infrastructure Distributing Biology:

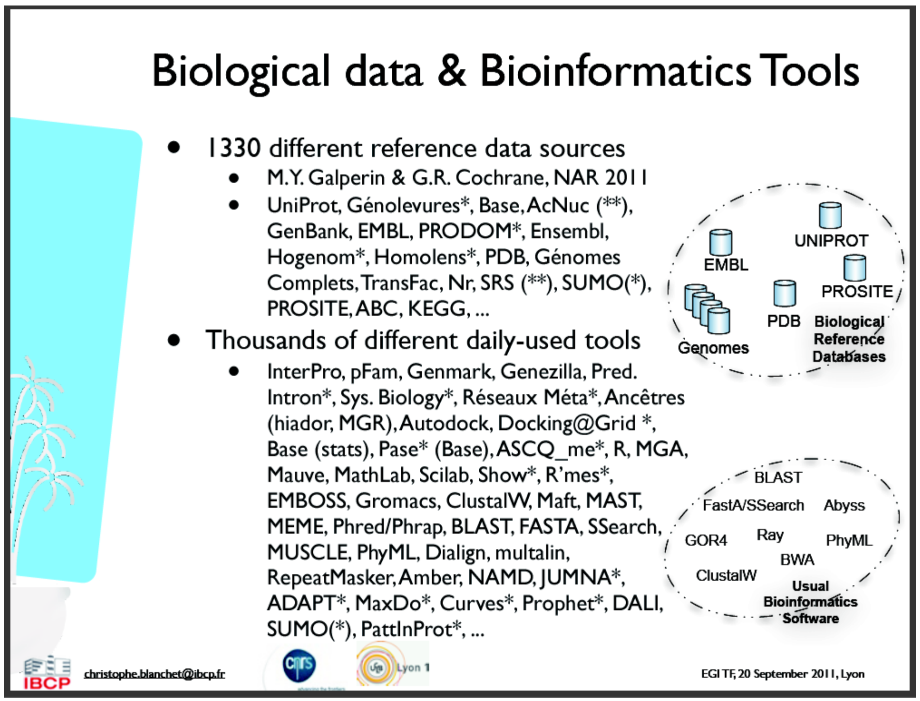

- Big task to port them all to your favorite architecture
- Is true for many scientific fields
- HPC Cloud:
  - Provides a different solution to “porting applications”
  - Provides the possibility to scale up your Desktop environment to HPC

40 JAAR 1971 2011 sara

The SARA / BiG Grid Cloud project

BiG Grid the dutch e-science grid

- The beginning
  - Started in 2009 from a powerpoint presentation
- Pilot
  - 3 months with 10 users on 48 core cluster
    - Participation was a “price” in a competition
    - Very succesfull
- Preparing for Operations
  - 2010 BiG Grid grants funding for project
    - 2010 – 2011 Product development
    - 50 usergroups on 128 core cluster
  - 2011 Big Grid grants funding for HPC hardware
- Now
  - Presenting a production infrastructure
  - HPC Cloud day 4th of October (slides and video is online at www.sara.nl)

Intro: a search for Use Cases

SARA has initiated a small scale experiment to investigate the use of Cloud Computing. In order to evaluate and share the experiences with the e-science community we are seeking researchers and developers with specific application needs, who want to make use of this opportunity to experiment with state of the art technology. We are especially interested in applications which are difficult or near impossible to run on our existing platforms (Huygens, Lisa, Grid), but do run in your local environment.

We invite you to send in your applications and ideas from which a selection will be made. Sara will support the implementation of working demonstrations as proof of concept on our test Cloud. Background, deadlines and procedure are described below.

Why Cloud Computing ?

Cloud computing is a broad concept, tightly related to virtualisation of resources. Using virtualised resources, a dynamic, scalable and more flexible set of computational services can be obtained. End-users or developers can create, select or configure their own operating system images and thus completely configure the cluster resources, this includes images running on multiple cores and possibly with shared memory (OpenMP/MPI).

Flexibility is obtained through the freedom in choice of operating system and thus possible applications to use. The time spent in porting applications to an existing HPC platform can be significantly reduced with Cloud Computing, because Cloud Computing gives end-users the ability to duplicate the environment where their application is currently running.

Virtualisation does come at a cost. However, some benchmarks on modern CPUs show the overhead of a VM compared to running jobs directly on the same hardware to be only between 2 and 5%, which is negligible. http://www.grid.phys.uvic.ca/documents/reports/files/chep09_vmbenchmark.pdf

HPC Cloud Project Publicity & dissemination

- In 2009 we were the first in the global Grid community to offer this to end users
- We started international interest group and presented at many events (ECEE, EGI, Sienna Cloud Scape, etc) and institutions
- We are in the top 5 search results for “HPC Cloud” and “HPC Cloud Computing” at Google
- News articles on paper and in online publications (HPCCloud, HPCWire, Computable, etc)

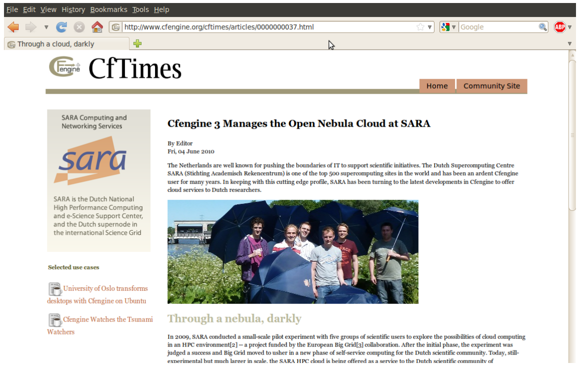

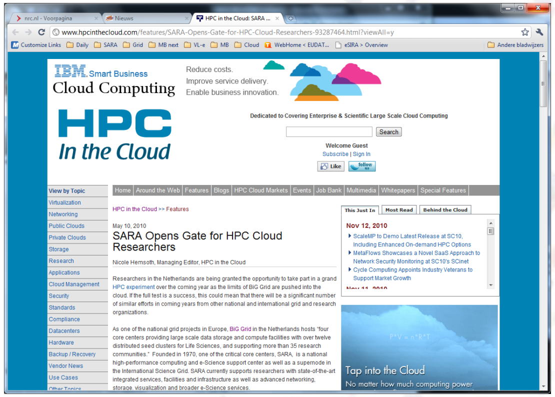

User participation
30 involved in Beta testing

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	Estimate an execution time of existing bioinformatics 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
30 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				al Science
4	A user friendly cloud-based inverse modelling environment				o-ecology,
5	Real life HPC cloud computing experiences for MicroArray analyses				artment, rmatics Unit,
6	Customized pipelines for the processing of MRI brain data				ng Group, mus MC
7	Cloud computing for historical map collections: access and georeferencing				raphy, UvA
8	Parallellization of MT3DMS for modeling contaminant transport at large scale				
9	An imputation pipeline on Grid Gain				formatics f groningen
10	Regional Atmospheric Soaring Prediction				o-ecology,
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

Field	# projects
Bioinformatics	14
Ecology	4
Geography	3
Computer science	4
Alpha/Gamma	5

Virtual HPC Cluster

- A true HPC cluster with dedicated storage, network and compute nodes
- User has full control
- **Upload** and run **any** image
- Custom environment examples
  - A much bigger workstation and/or many identical workstations
  - Closed source software
- Different environments in same cluster
  - Example: combine Linux with windows nodes
- High mem, disk, I/O, CPU and in any combination
- Long running jobs

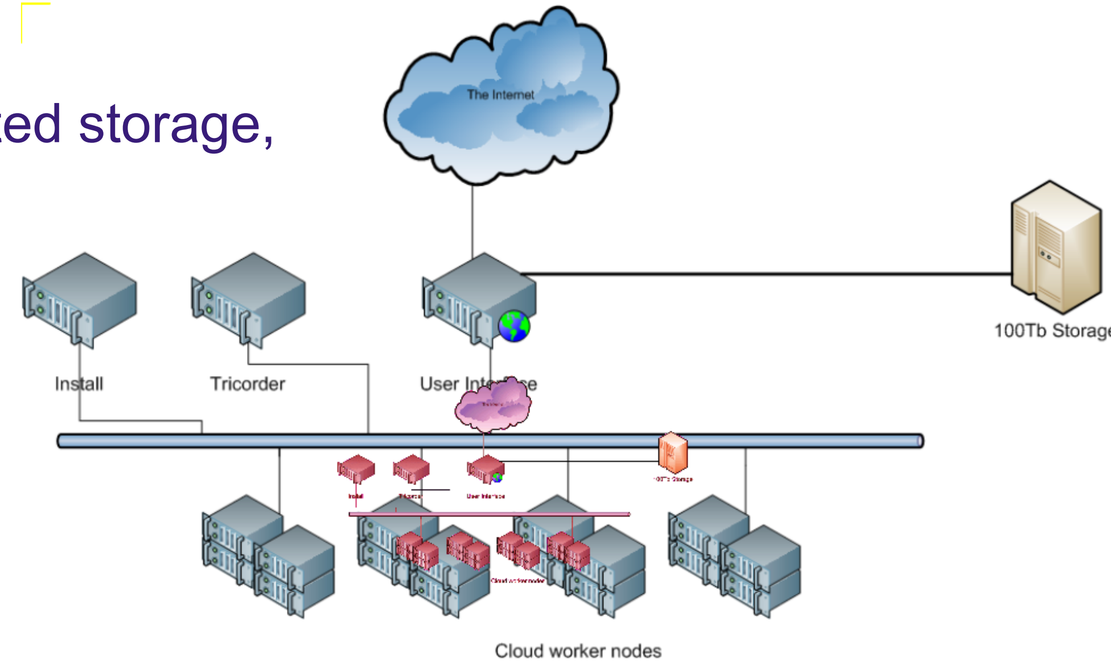

HPC Cloud Application types

- Applications with different requirements can co-exist on the same physical host and in the same virtual cluster
- Single node (remote desktop on HPC node)
- Master with workers (standard cluster)
- Token server
- Pipelines/workflows
  - example: MSWindows+Linux
- 24/7 Services that start workers

HPC 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

HPC Cloud Architecture

Calligo *“I make clouds”*

40 JAAR 1971 2011 sara

BiG Grid the dutch e-science grid

1 Node:

- CPU Intel 2.13 GHz 32 cores (Xeon-E7 "Westmere-EX")
- RAM 256 Gbyte
- "Local disk" 10 Tbyte
- Ethernet 4*10GE

Total System (19 nodes)

- 608 cores
- RAM 4,75TB
- 144 ports 10GE, 1-hop, non-blocking interconnect (fabric: 2.5 Terabit/s full-duplex,4.5usec port to port latency
- 400TB shared storage (ISCSI,NFS,CIFS,WebDAV...)
- 11.5K specints / 5TFlops

- Novel architecture
  - Low latecy 10GE switch chassis at the heart, with mountable storage
  - Per node High mem, high core, 4* network (but smaller partitions possible)
  - Fast and dynamic changes in configurations and partitions
  - Security and safety through strict separation of users + monitoring network
- Invested in Expandability
  - Network switch and storage controllers can be expanded to support more compute nodes
  - Future architecture focus might change, depending on usage patterns
- Virtualization overhead <10%
  - Very flexible for any application. “Just great” is often good enough

Platform and tools:
Redmine collaboration portal
Custom GUI (Open Source)
Open Nebula + custom add-ons
KVM + Libvirt

intel DELL DataDirect NETWORKS ARISTA clustervision

Cloud = Expensive???

- Please Note:
  - The hardware is what makes it "High performance computing"
    - We choose this type of hardware to fill a specific niche in our offerings
  - The Software stack is what makes it "Cloud"
  - You can do this with any type of hardware (but be careful with Infiniband)!!
  - Actually we give much less user support
    - expensive hardware but cheaper in personnel

HPC Cloud Architecture vs other HPC Architectures

System	Node	Characteristics
Huygens National Super 	CPU Power6, 4.7Ghz, 32/64 cores RAM 128/256 GB "local disk" 8Tbyte Infiniband 8*20 Gbit/s	Network, storage, and compute are **fully balanced** resulting in a **highly optimized** system for I/O- intensive, parallel applications
LISA National Compute Cluster 	CPU Intel 2.26Ghz 8/12 cores RAM 24 Gbyte Local disk 65Gbyte/200Gbyte Infiniband 20Gbit/s	Network, storage, and compute are **optimized** for **efficient** parallel applications
Grid 	CPU Intel 2.2 Ghz 8 cores RAM 24 Gbyte Local disk 300Gbyte Ethernet 1 Gbit/s	Network, storage, and compute are **optimized** for **high throughput** of data and single core data processing
Cloud 	CPU Intel 2.13 GHz 32 cores (Xeon-E7 "Westmere-EX") RAM 256 Gbyte "Local disk" 10 Tbyte Ethernet 4*10GE	Network, storage, and compute are **dynamically scalable and configurable** for both large and small workloads

Cloud security & trust

Chief principle: !!! CONTAINMENT !!!

- Protect
  - the outside from the cloud users
  - the cloud users from the outside
  - the cloud users from each other
  - Not possible to protect the cloud user from himself because user has full access/control/responsibility
    - ex. virus research must be possible
- Keep each VM inside its own sandbox, completely separated from others and its host
- Private Vlan per user
  - Access to outside only through proxy
  - Access from outside only through VNC on proxy (port forwarding)
- Public ip is allowed (both incoming and outgoing), but Source & Target (range) need to be specified
  - Self service & Dynamically: rules are autogranted
  - Traffic is monitored
  - Open Ports are monitored
- Specialized hardware helps, but not necessary
  - QinQ, Bridges, IP tables

Calligo Network

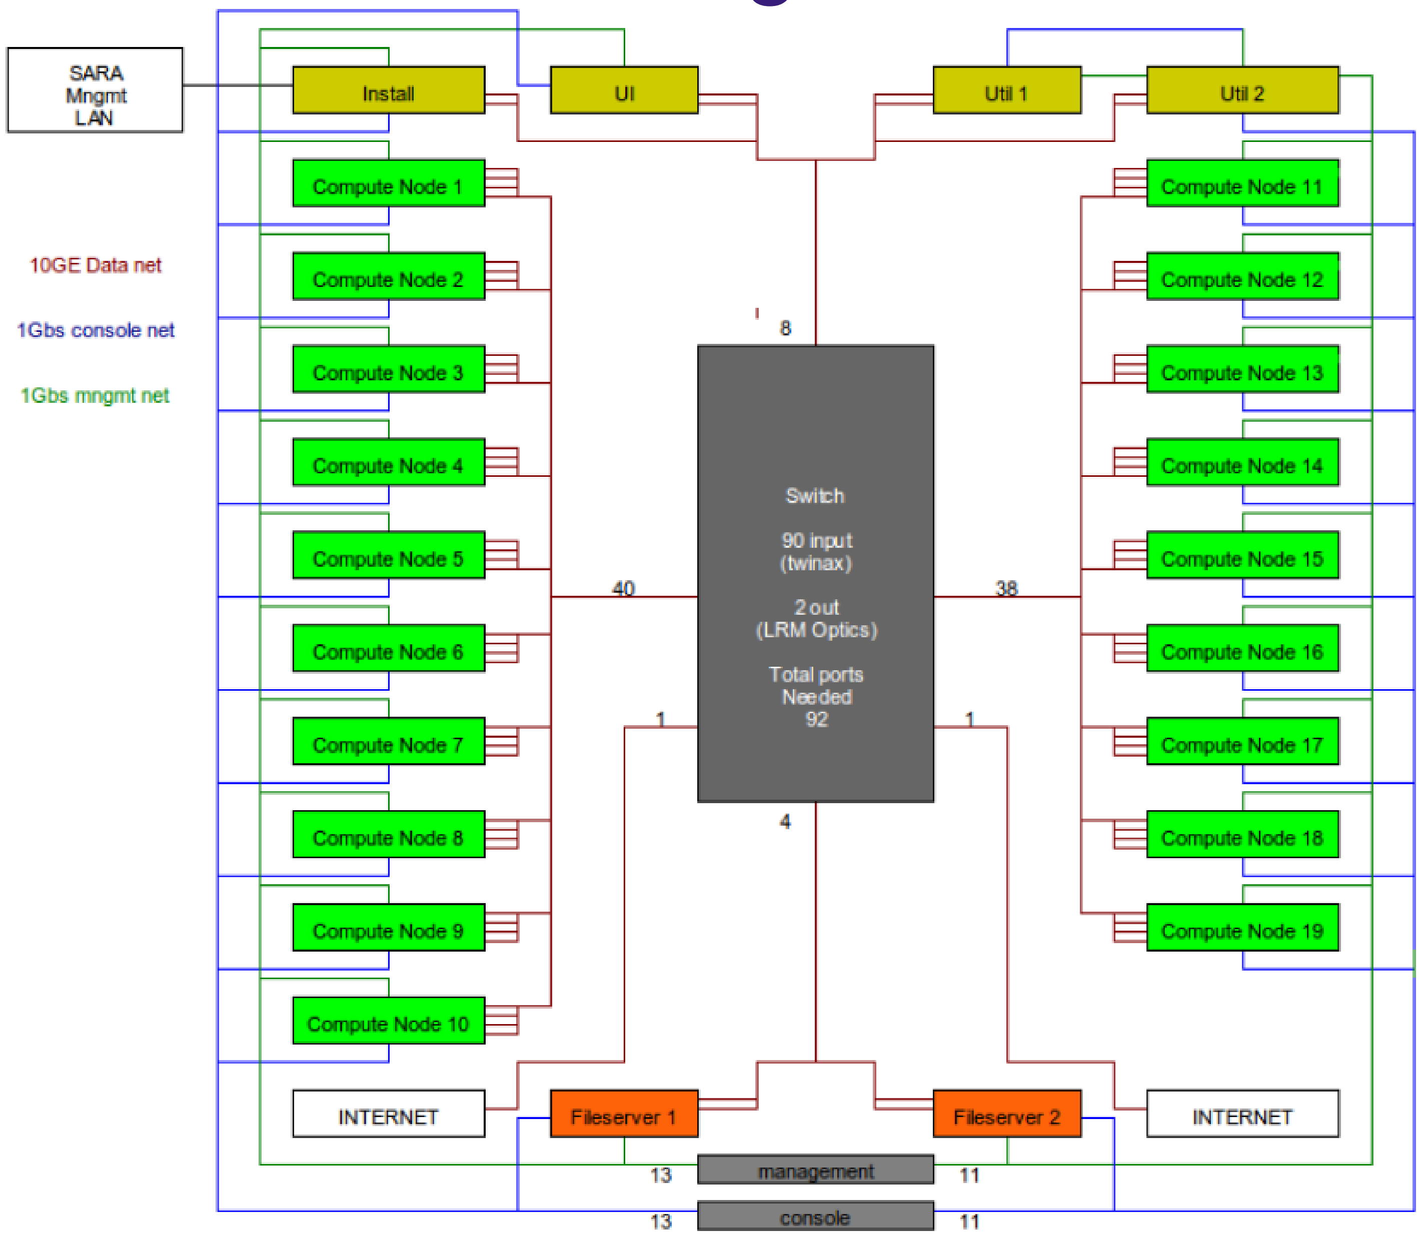

- 1 entry point from administration network
- No connection from admin to user space
- Almost full redundancy
- Switch fabric is impressive: 3.75 Terabit full-duplex with 4.5usec port to port latency
- Dedicated lightpaths to various systems in the Netherlands

Cloud Security
"Public" side

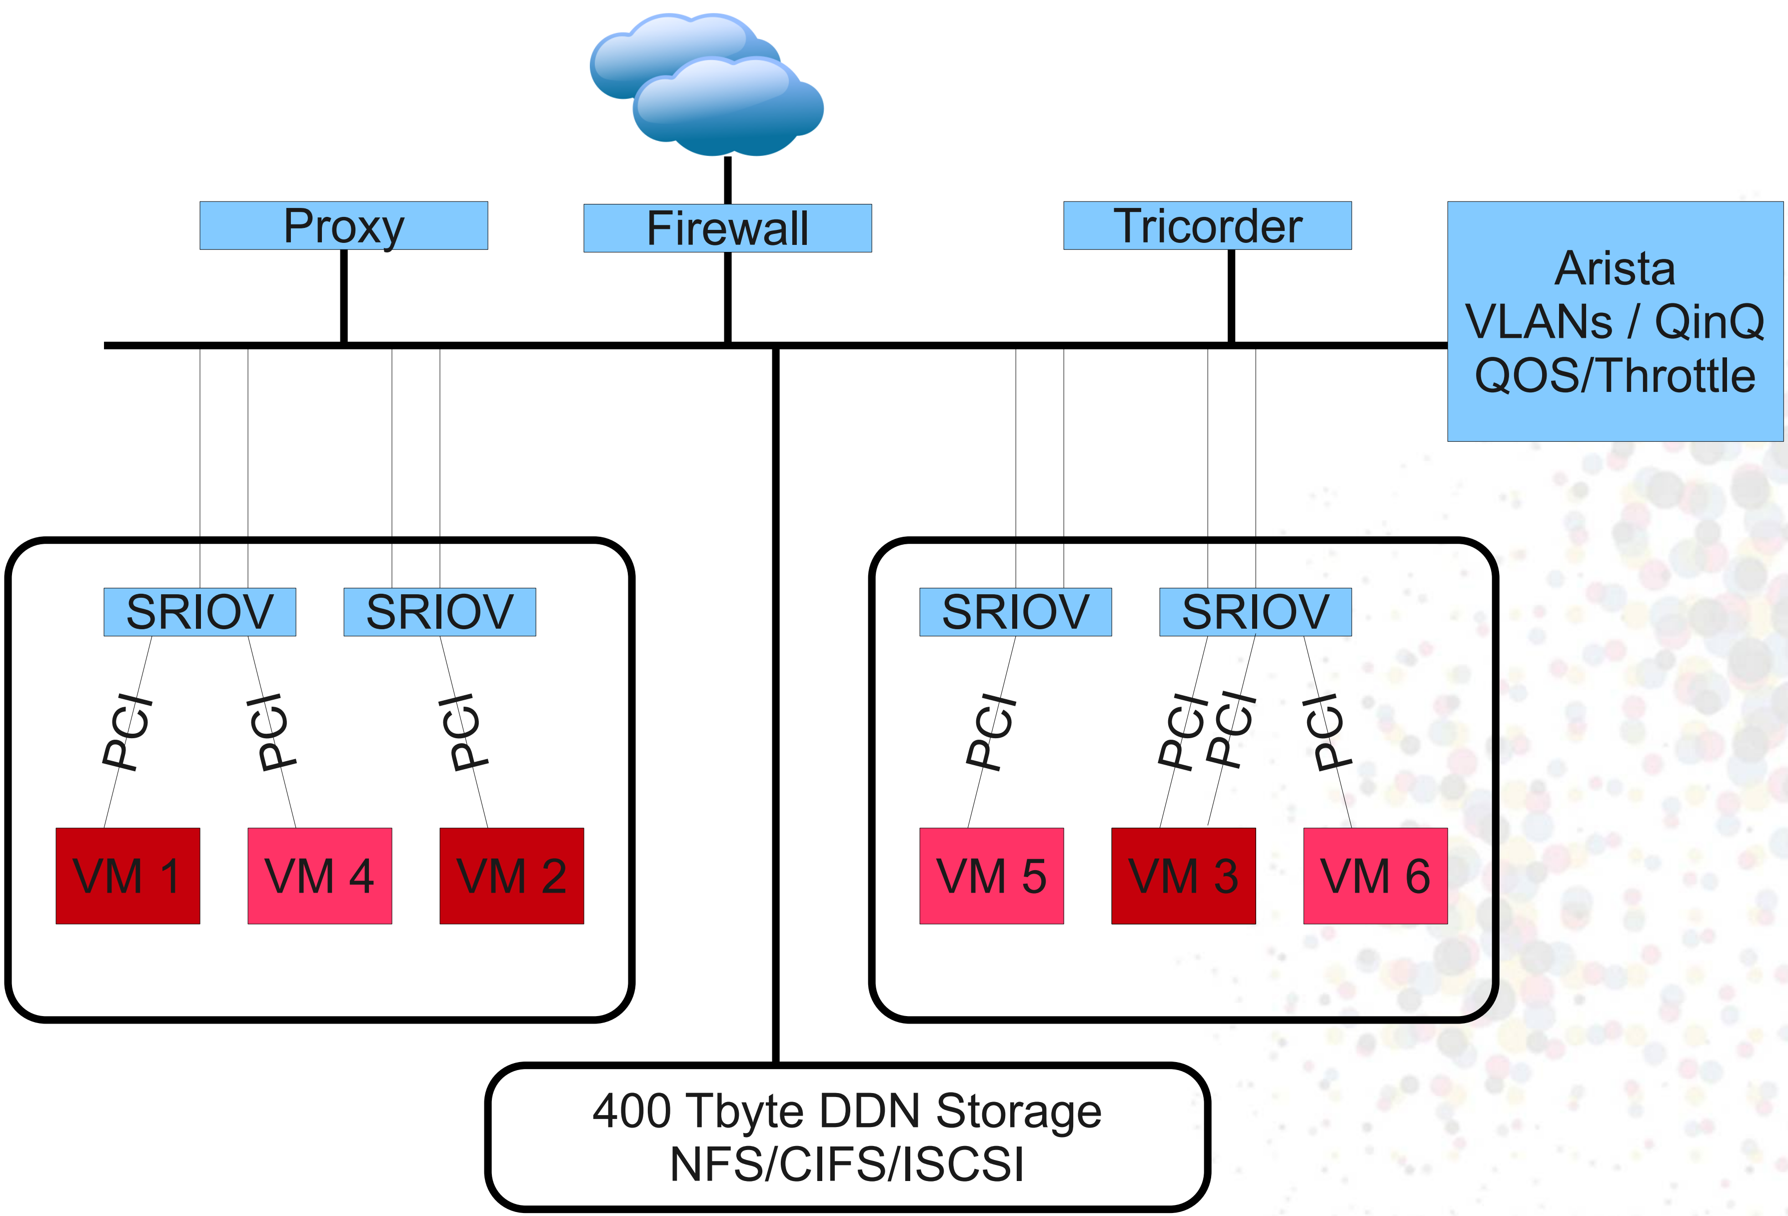

Cloud Security (VLANS + monitoring)

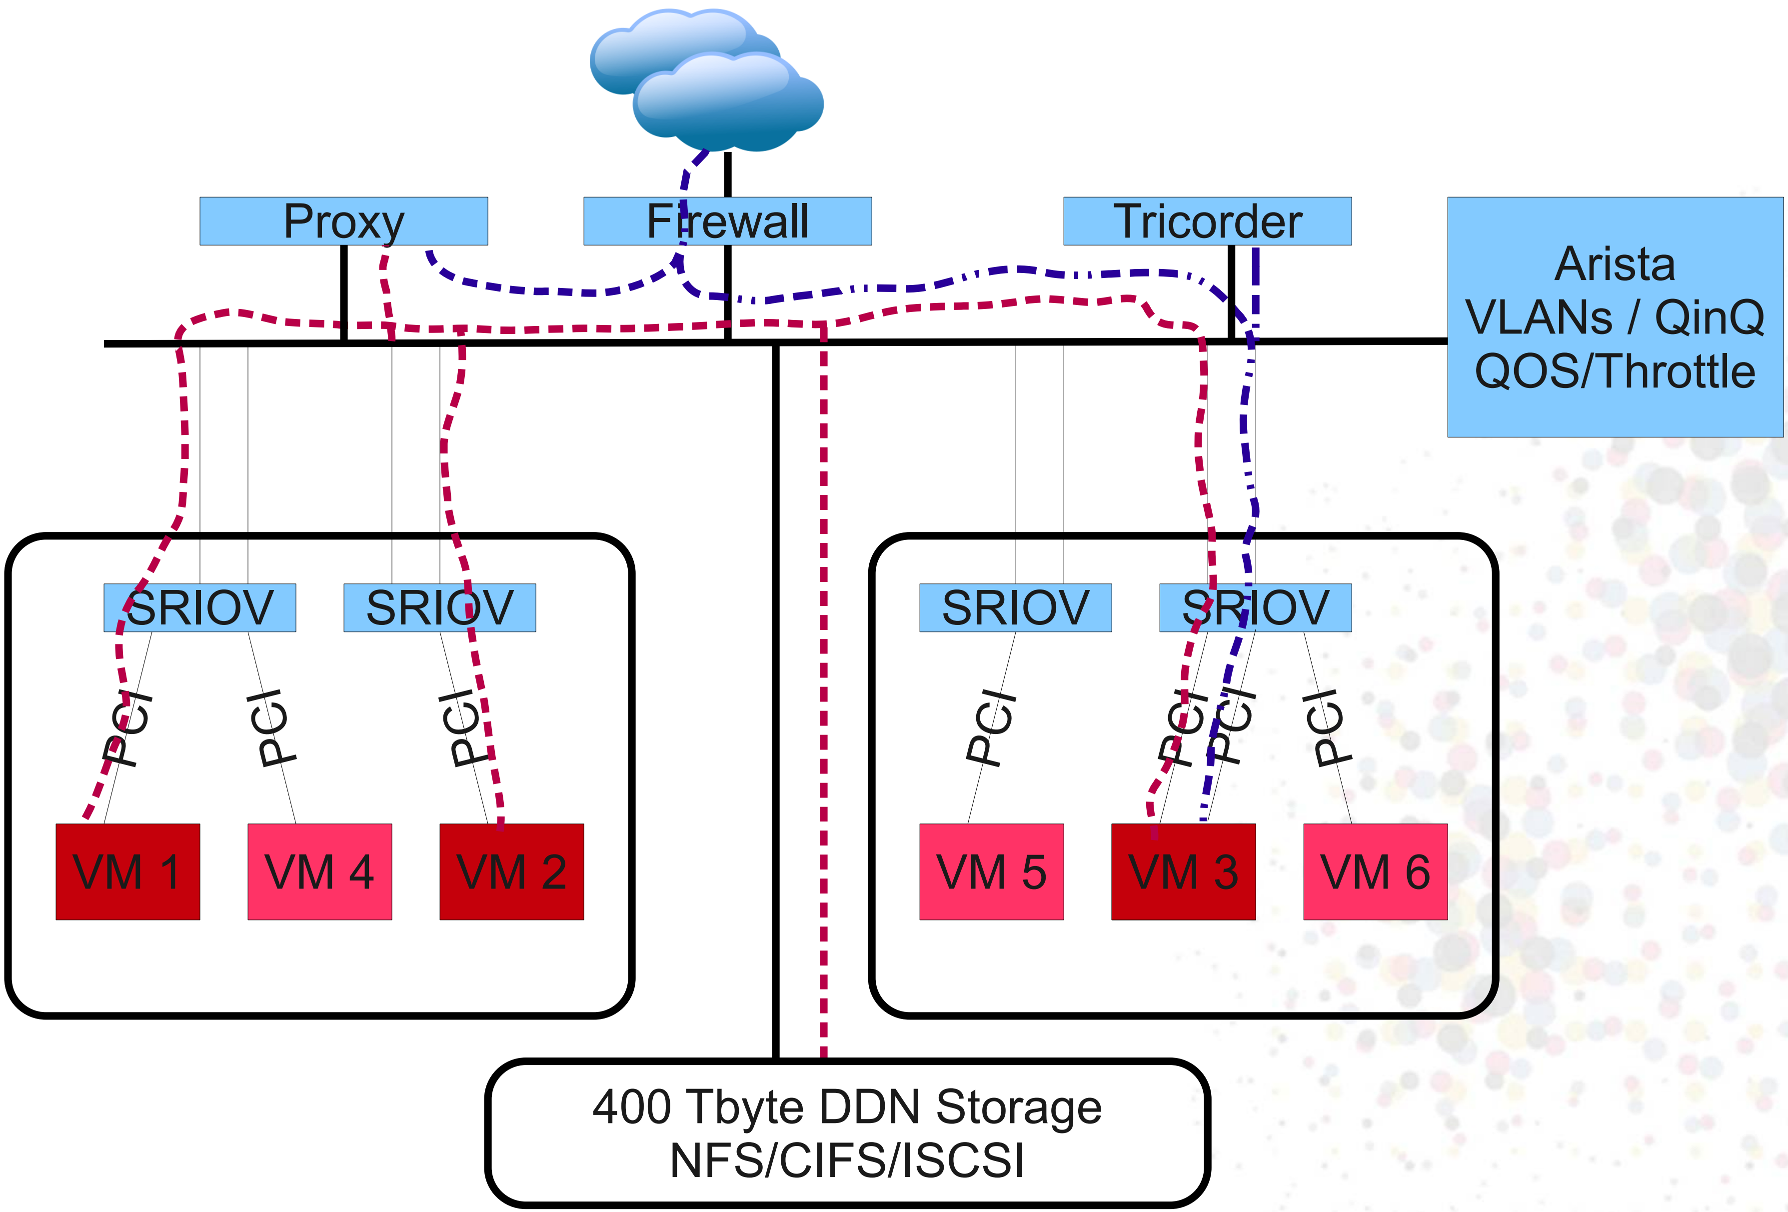

GUI + Collaboration Portal

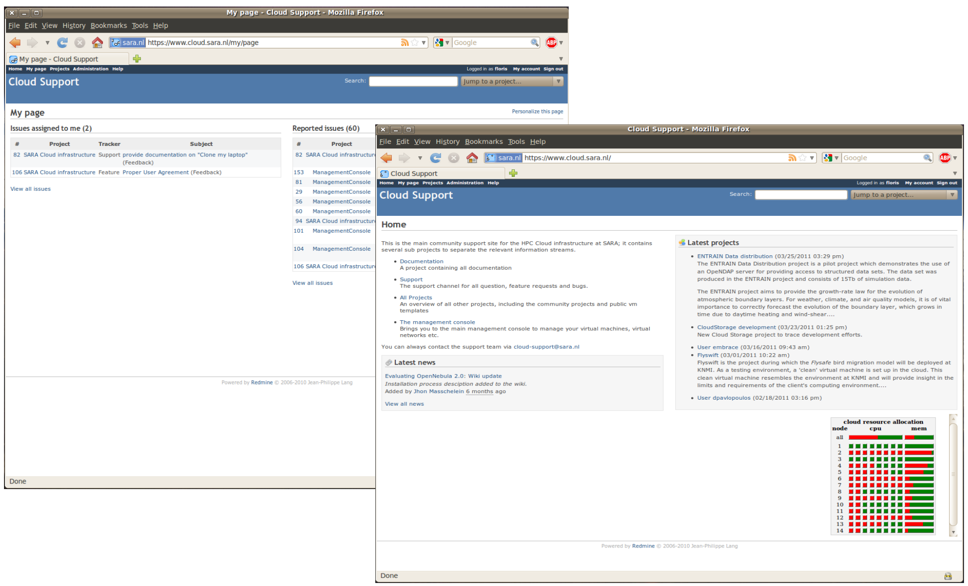

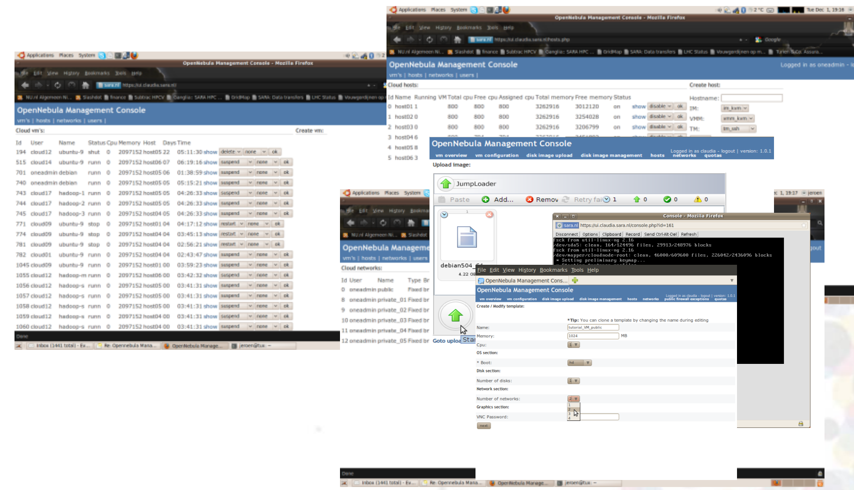

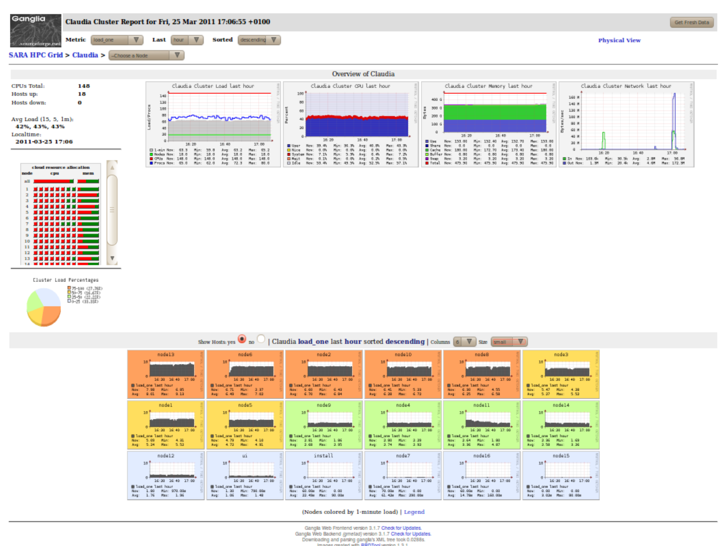

For current and future users

- Most current users only asked a few support questions when they start to use it, it is easy to do support on this system ;-)
- In the Pilot and Beta phase, the hardware was not ideal for this type of system and applications. The new hardware is/will be much bigger and better!
- On the HPC Cloud: Do not change your HPC applications for the system, but rather adapt the environment they run in
  - The result will be: a shorter time to a solution

Thank you!

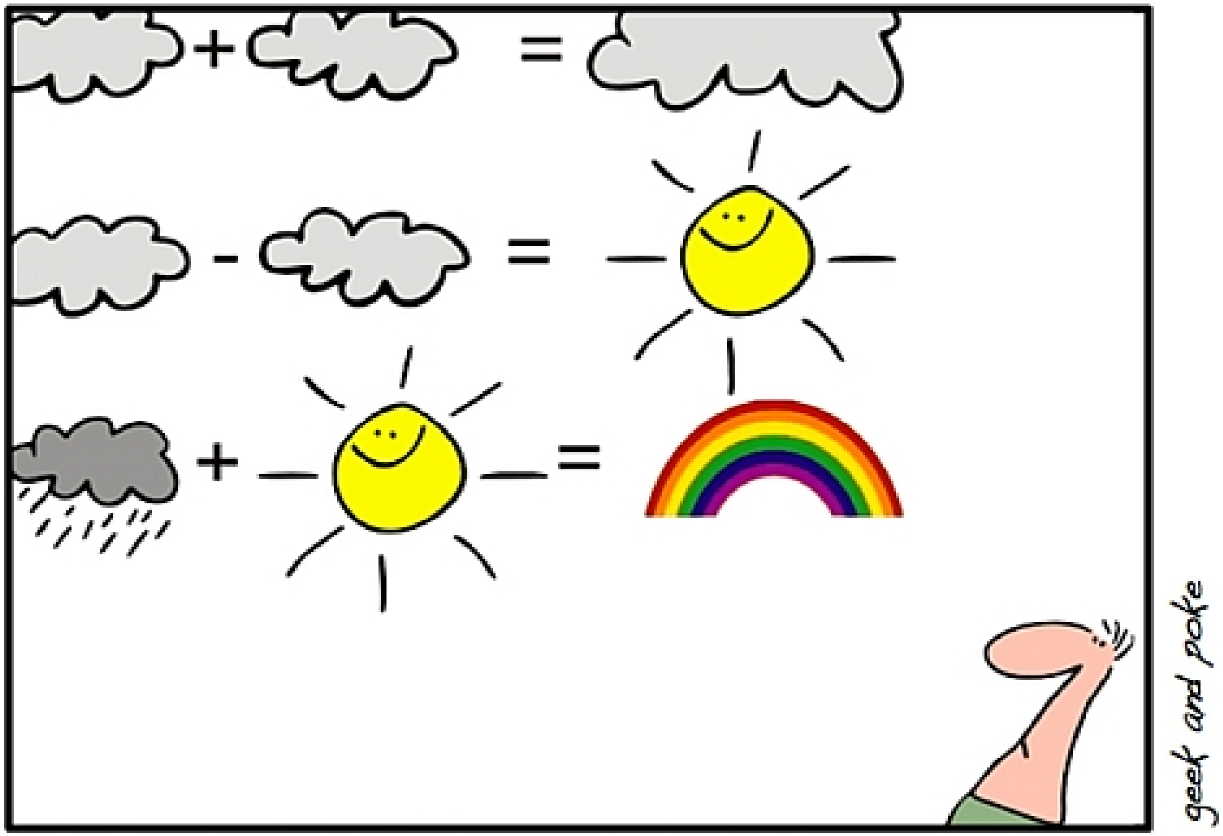

SIMPLY EXPLAINED – PART 17:
CLOUD COMPUTING

Questions?

www.cloud.sara.nl

Photo's: http://cloudappreciationsociety.org/