Karlsruhe Institute of Technology

Clouds and Grids

DESY – Hamburg – July 6th 2009

Christian Baun

Forschungszentrum Karlsruhe
in der Helmholtz-Gemeinschaft

Universität Karlsruhe (TH)
Forschungsuniversität • gegründet 1825

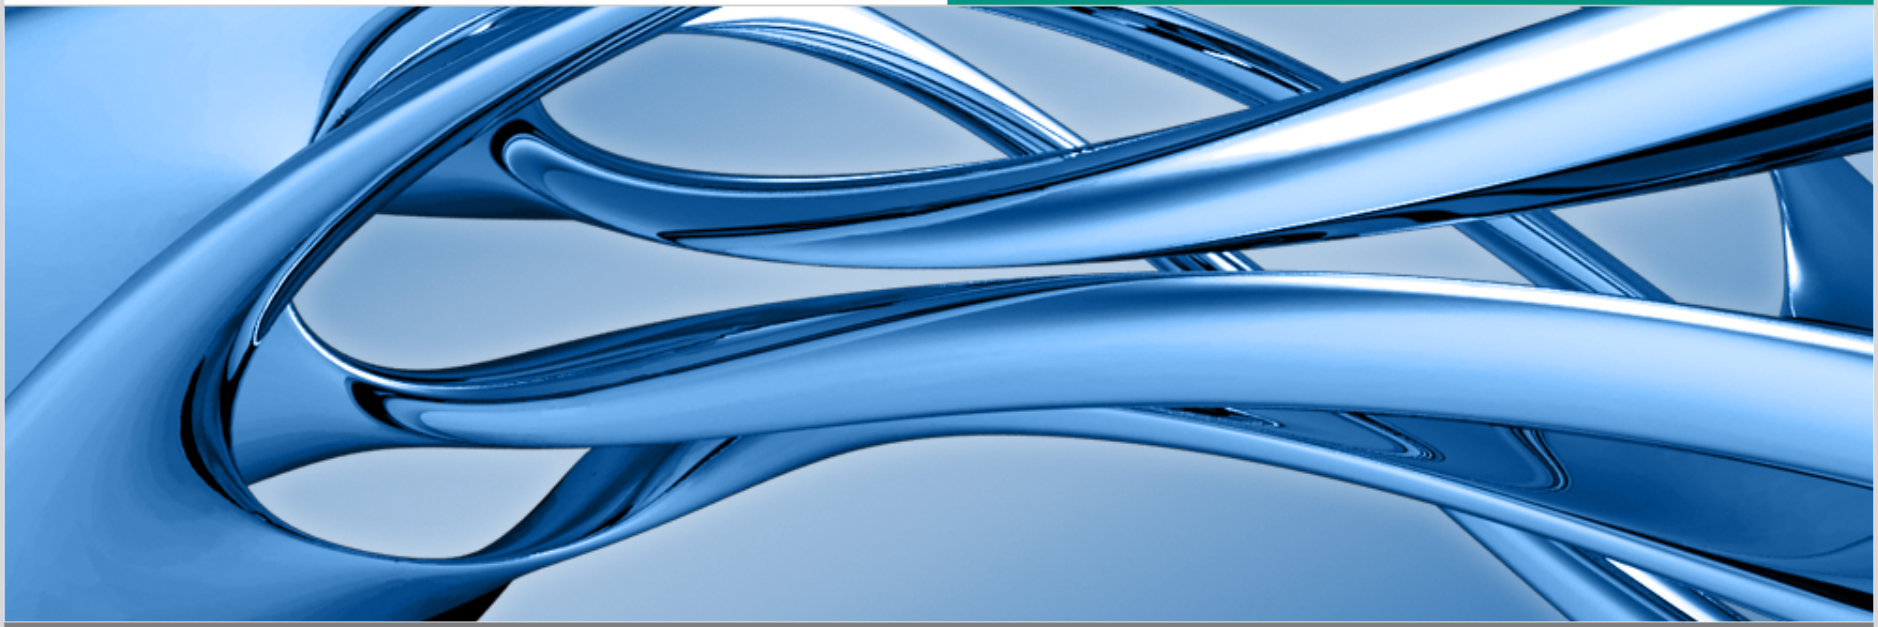

Cloud Computing - big topic/hype

- Big topic (hype) in the press since autumn 2008

Gartner Identifies Top Ten Disruptive Technologies for 2008 to 2012

MELBOURNE, Australia, May 28, 2008 — Social networking technologies, web mashups, multicore and hybrid processors and cloud computing are amongst the ten most disruptive technologies[1] that will shape the information technology (IT) landscape over the next five years, according to research and advisory firm Gartner, Inc.

QUOTABLE

Larry Ellison on cloud computing buzzword: "Complete gibberish"

By Jackson West, 1:20 PM on Fri Sep 26 2008, 1,766 views

Cloud computing is a trap, warns GNU founder Richard Stallman

Web-based programs like Google's Gmail will force people to buy into locked, proprietary systems that will cost more and more over time, according to the free software campaigner

Bobbie Johnson, technology correspondent
guardian.co.uk, Monday 29 September 2008 14.11 BST

Cloud computing: Hot technology for 2009

Proceed with caution

By Neal Weinberg , Network World , 01/05/2009

As we arrive at 2009, cloud computing is the technology creating the most buzz. Cloud technology is in its infancy, however, and enterprises would be wise to limit their efforts to small, targeted projects until the technology matures and vendors address a variety of potentially deal-breaking problems.

Cloud Computing: more facts please!

- Good definitions of Cloud Computing are rare

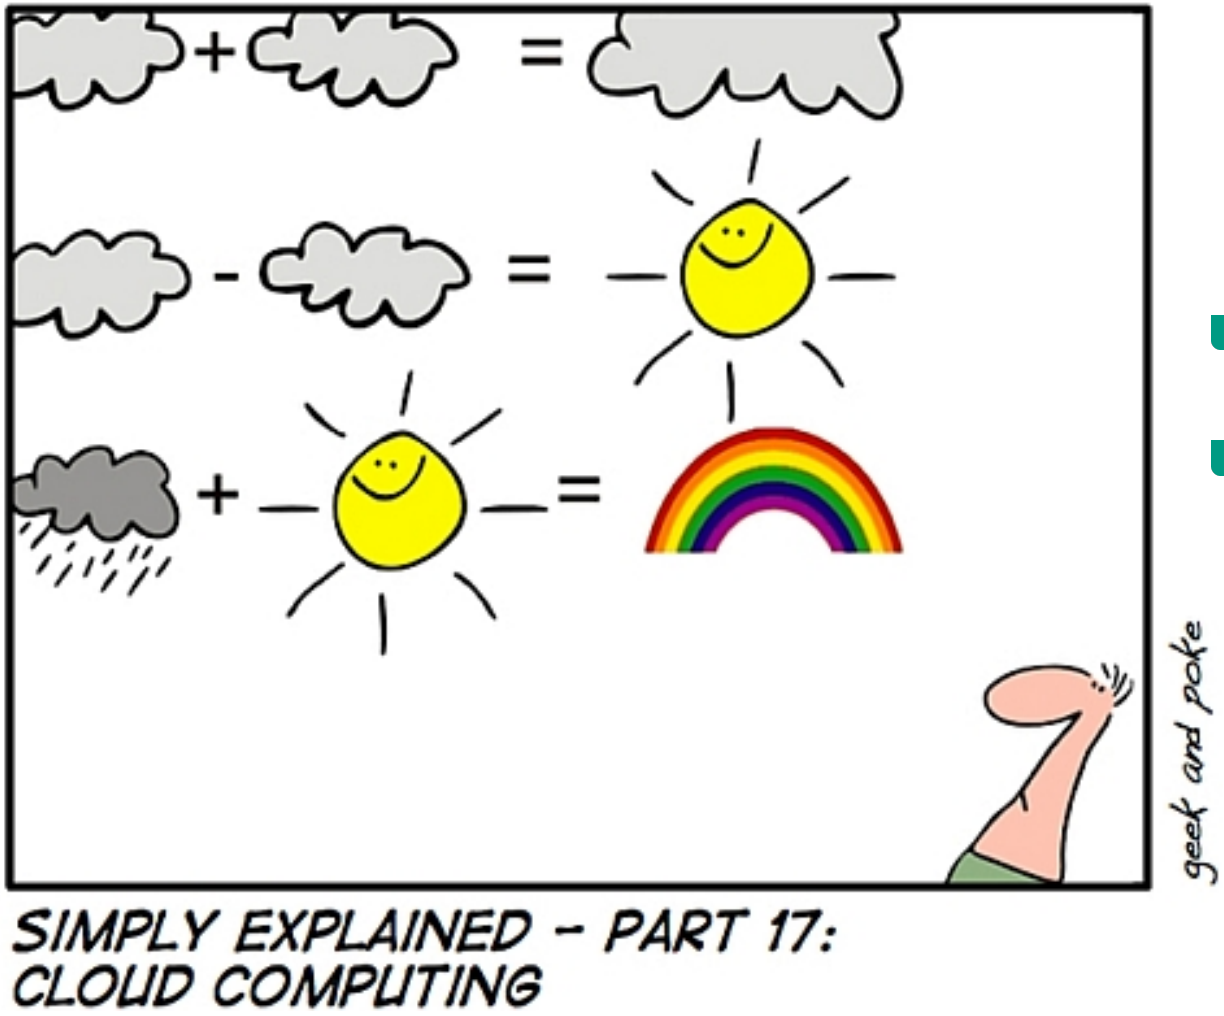

- What is behind Cloud Computing?
- Why do we need Cloud Computing?

Why do we want/need Cloud Computing?

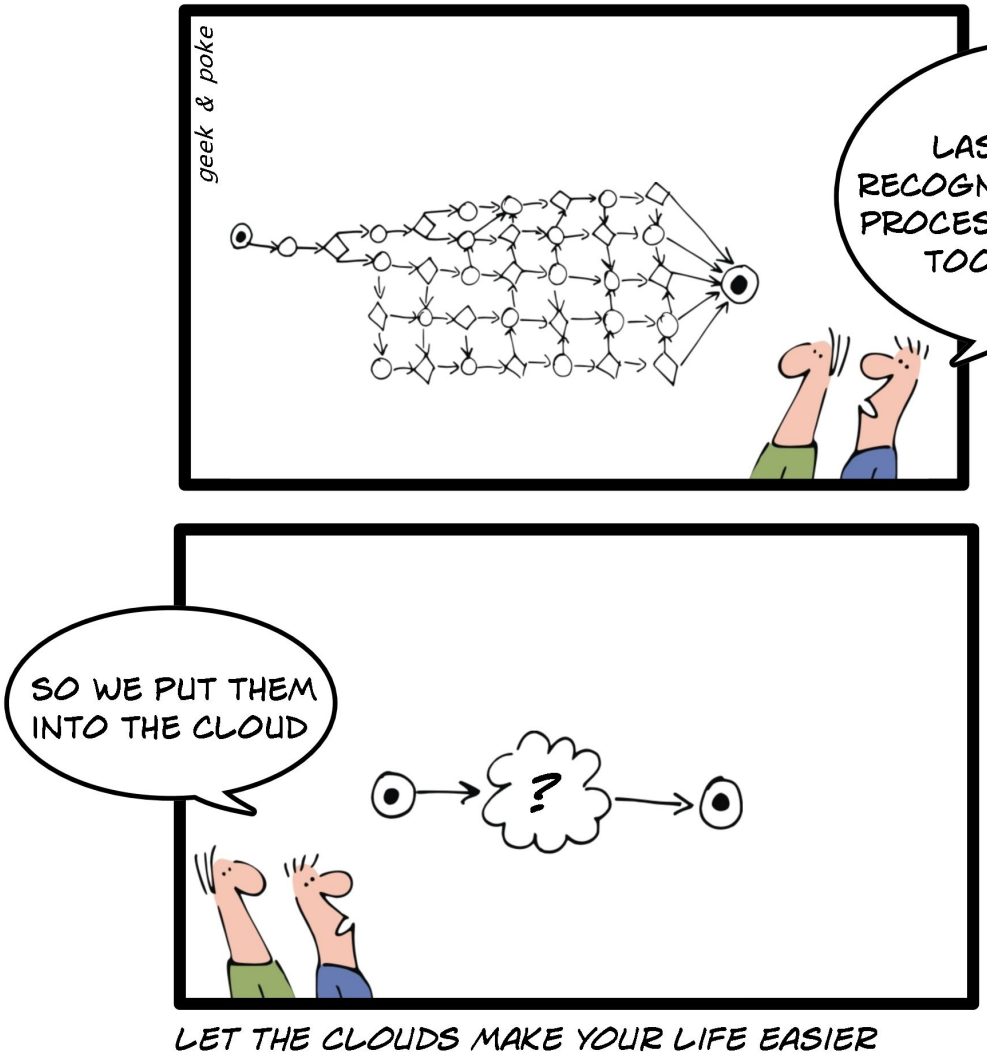

- Cloud Computing shall:
  - revolutionize IT
  - reduce (eliminate) complexity
  - getting IT more flexible (elastic)
  - reduce cost
  - provide easy resources access
  - fulfill users demands
  - emancipate the users
  - ...
- Remember: GRID !?!?!?

Definitions (Cloud / Grid)

- **Cloud Computing is on-demand access to virtualized IT resources that are sourced inside or outside of a data center, scalable, shared by others, simple to use, paid for via subscription or as you go and accessible over the web.**
  Dr. Behrend Freese (Zimory GmbH)
- **A computing Cloud is a set of network enabled on demand IT services, scalable and QoS guaranteed, which could be accessed in a simple and pervasive way.**
  Dr. Marcel Kunze (SCC/KIT)

- **Grid computing is coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organizations.**
  Ian Foster (Argonne National Laboratory)
- **A computational grid is a hardware and software infrastructure that provides dependable, consistent, pervasive, and inexpensive access to high-end computational capabilities.**
  Ian Foster & Carl Kesselman

Cloud Comuting = Grid Computing ?!

Cloud Computing	Grid Computing
One or few data centers containing heterogeneous or homogeneous resources under central control	Geographically distributed, heterogeneous resources without central control. Follows the principle of virtual organizations
Virtualized resources	Physical resources
Fully automated Services (Industrialization of IT)	Handcrafted Services (Mostly Manufacture)
Commercial business model (Pay-as-you-go)	Publicly funded. Usage is for free if the resource owner approves resource access
Easy to use and deploy. No complex user interface required	Difficult to use and deploy
Used primarily in industry and business (Big chance for Startups)	Used first and foremost by research and academic organizations (Large-scale scientific projects like LHC)

THE Grid-Project: LHC Computing Grid (1)

http://lcg.web.cern.ch/LCG/

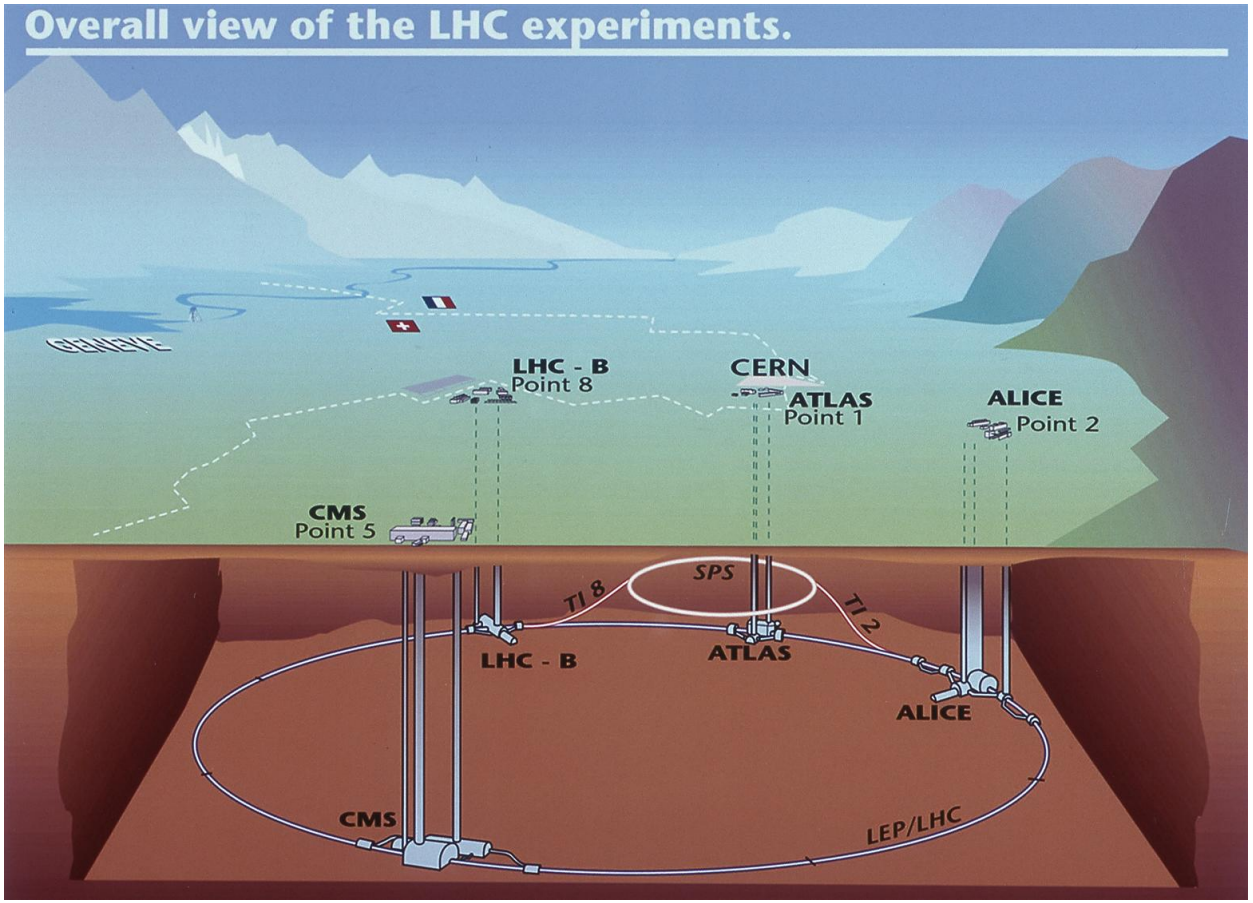

Source: http://guillegg.wordpress.com

- The LHC is the worlds largest and highest-energy particle accelerator
- 4 main detectors
- Located at CERN. The European Organization for Nuclear Research
- Used to test various predictions of high-energy physics
- Big goal: confirm or refute the existence of Higgs boson („God particle“)

THE Grid-Project: LHC Computing Grid (2)

http://lcg.web.cern.ch/LCG/

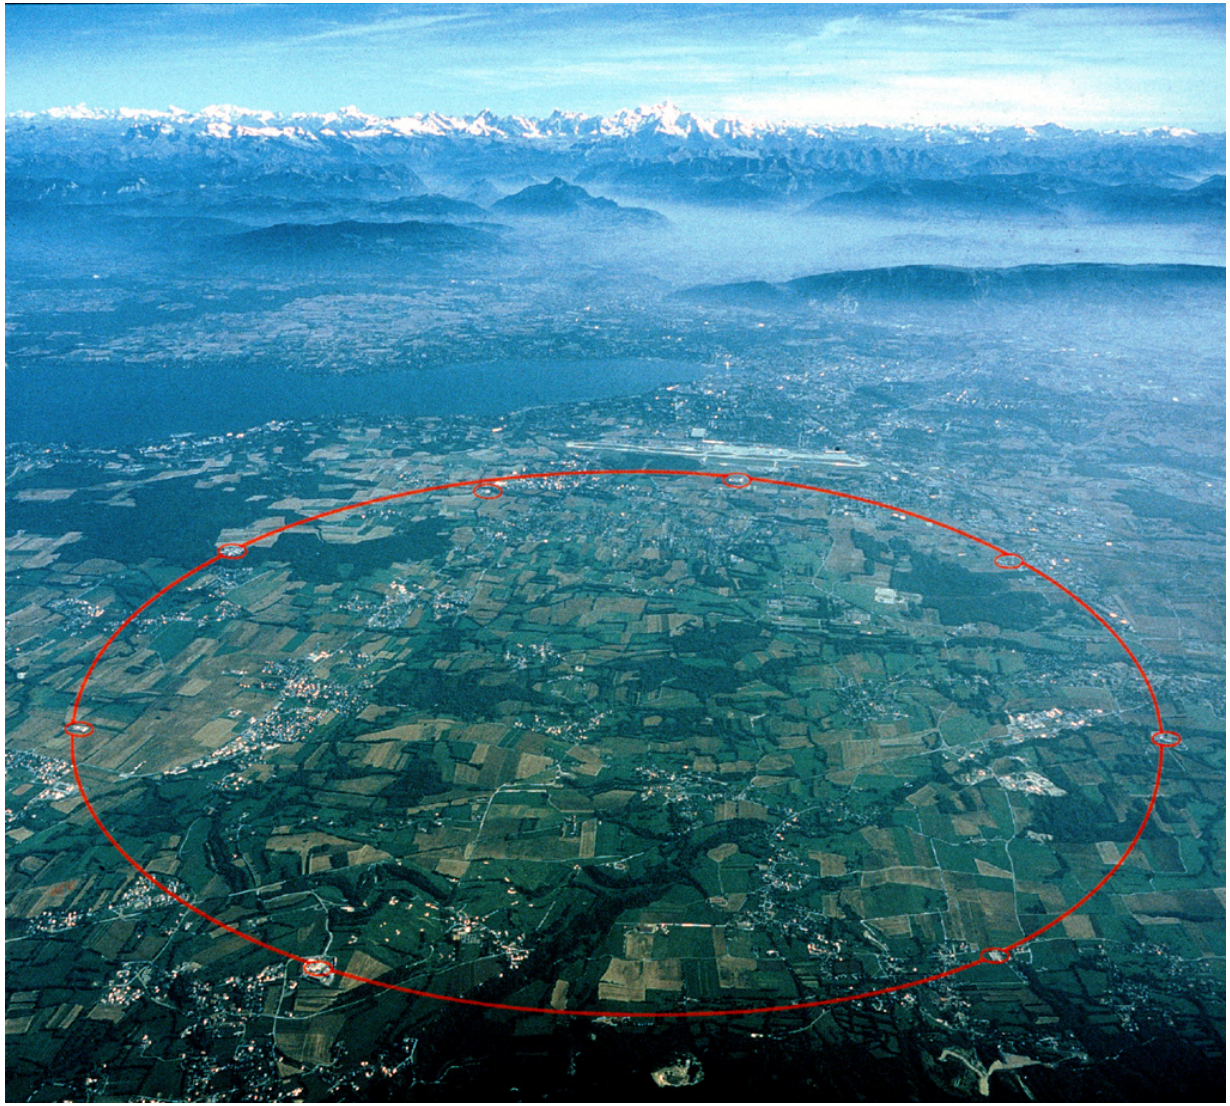

Source: http://www.nssp.uni-saarland.de

- The LHC generates 10 - 40 Petabyte of data for all experiments per year
- Data is stored and analyzed distributed inside the LHC-Grid
  - Largest computing grid in the world
  - > 150 Sites
  - > 20000 Servers

Three technical Types of Cloud Services

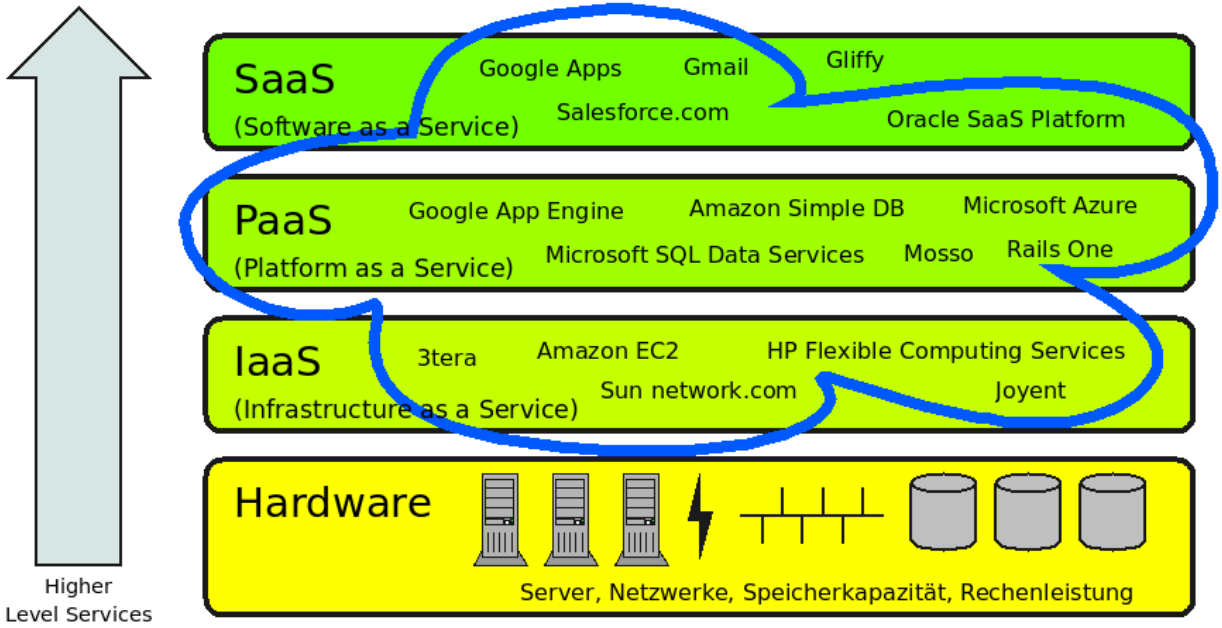

- **SaaS**
  - **Provides enterprise quality software is run by a service provider**
  - **The user has no need to worry about installation, administration or updates**
- **PaaS**
  - **Appears as one single large computer and makes it simple to scale from a single server to many**
  - **No need to worry about the operating system or other foundational software**
- **IaaS**
  - **Abstracts away the hardware (servers, network,…) and allows to run virtual instances of servers without ever touching a piece of the hardware**

Three organizational Types of Cloud Infrastrctures

- Public Clouds (respectively External Clouds)
  - Customer and Cloud provider do not belong to the same organization
  - Commercial business model (Pay-as-you-go)
- Private Clouds (respectively Internal Cloud or Intra Cloud)
  - Cloud services are provided from the users organization
  - Main reasons for Private Clouds: Security concerns and R&D
- Hybrid Clouds
  - Cloud services from one or more Public and Private Clouds are used
  - Use cases:
    - Public Clouds help to manage load peaks
    - Backup a Private Clouds data in Public Clouds

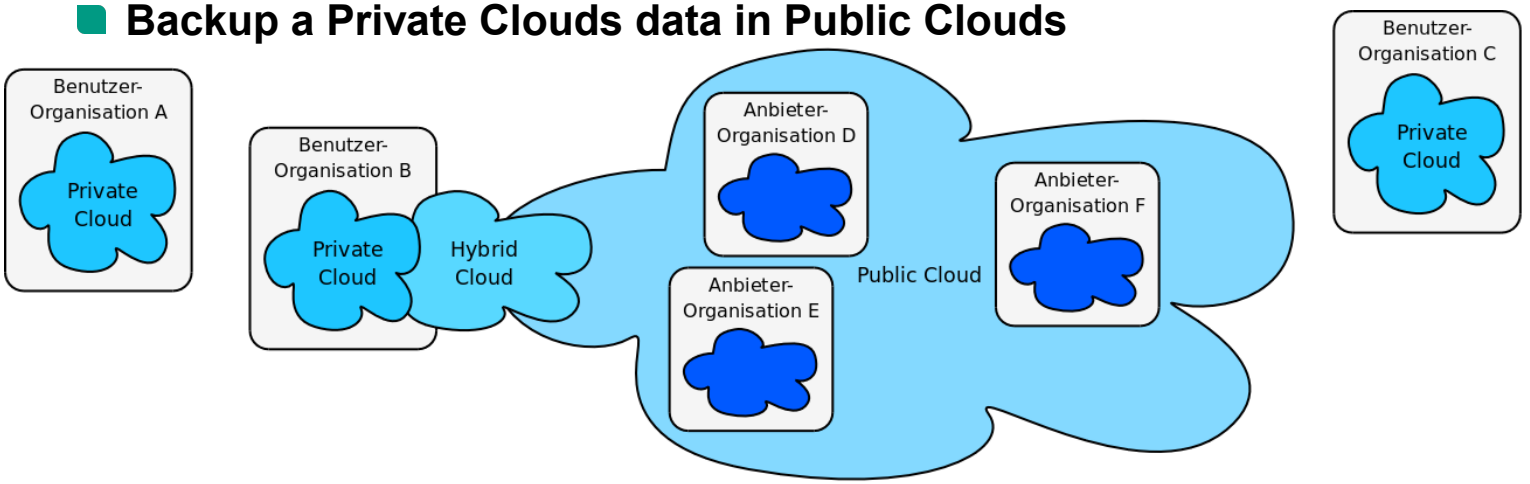

Commercial Cloud Offerings (Small Excerpt)

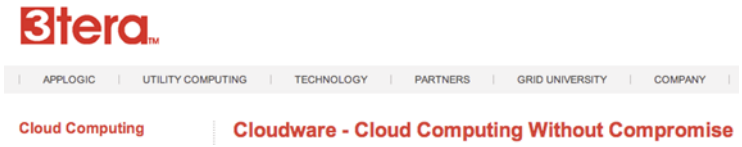

- **Problem: Commercial Cloud offers are usually proprietary and therefore not open for building Private Clouds and for Cloud systems research and development!**

Eucalyptus

http://eucalyptus.cs.ucsb.edu

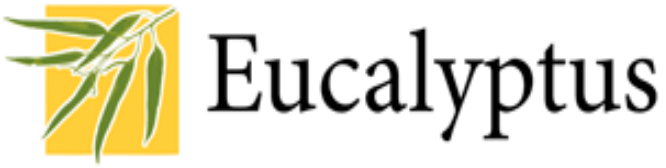

- EUCALYPTUS - Elastic Utility Computing Architecture for Linking Your Programs To Useful Systems
- Open Source software infrastructure for implementing Cloud Computing on clusters from UC Santa Barbara
- Developed at UC Santa Barbara
- Implements Infrastructure as a Service (IaaS)
- Gives the user the ability to run and control virtual machine instances (Xen, KVM) deployed across a variety of physical resources
- Interface compatible with Amazon EC2, S3 and EBS
- Includes „Walrus“, a storage service
- Potential to interact with the same tools, known to work with Amazon EC2, S3 and EBS
- Eucalyptus is an important step to establish an open Cloud Computing infrastructure standard

Eucalyptus - Components

http://eucalyptus.cs.ucsb.edu

Eucalyptus

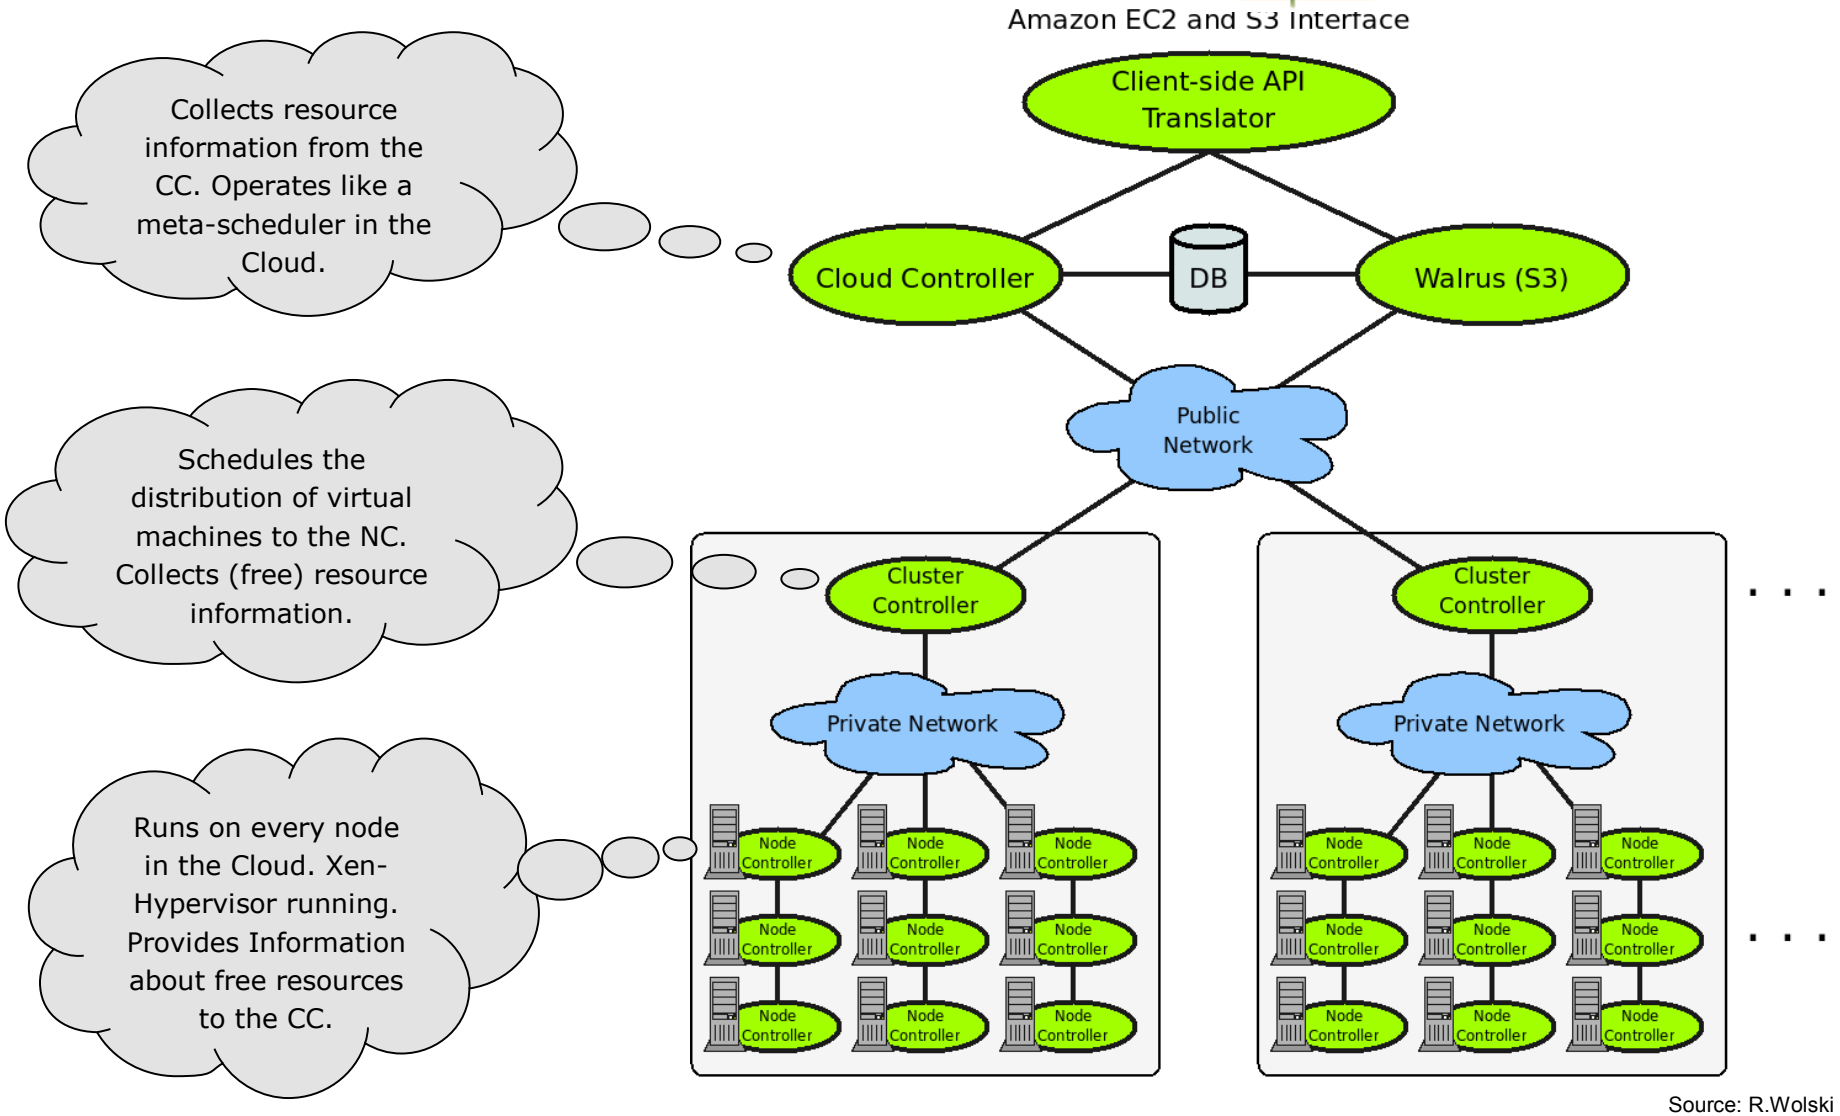

Source: R.Wolski

Eucalyptus at the FZK

- Cloud Installation I (Eucalyptus 1.4)
  - Runs stable
  - Plattform for performance testing
  - 2x IBM Blade LS20
    - 2 Single Core Opteron (2,4GHz)
    - 4GB RAM
  - 2x IBM Blade HS21
    - 2 Dual Core Xeon (2,33GHz)
    - 16GB RAM
- Cloud Installation II (Eucalyptus 1.5.1)
  - Under construction
  - 5x HP Blade ProLiant BL2x220c
  - 2 Server per blade:
    - 2x Intel Quad-Code Xeon (2,33GHz)
    - 16GB RAM

Eucalyptus (installation)

- **Binary packages exist for CentOS, openSUSE, Debian and Ubuntu**
  http://open.eucalyptus.com/downloads
  http://open.eucalyptus.com/wiki/EucalyptusAdministratorGuide_v1.5
- **Eucalyptus can be installed from source also**
  http://open.eucalyptus.com/wiki/EucalyptusSourceCodeInstallation_v1.5
- **Requirements: ≥ 1 computer running with Linux and a functional Xen Hypervisor or Kernel-based Virtual Machine (KVM)**
  - **For KVM, a *modern* CPU with AMD-V (Pacifica) or Intel VT (Vanderpool) is needed**
- **Amazon EC2 command line tools to control Eucalyptus**
  - **ec2-api-tools-1.3-30349**
  - **ec2-ami-tools-1.3-26357**

Use Eucalyptus

- **Shows the cluster's front-end hostname, free resources, instance types, available NCs:**

```
# ec2-describe-availability-zones verbose
AVAILABILITYZONE     Cluster1   iwrcgblade11
AVAILABILITYZONE     |- vm types      free / max    cpu    ram   disk
AVAILABILITYZONE     |- m1.small      0020 / 0024    1     128     10
AVAILABILITYZONE     |- c1.medium     0020 / 0024    1     256     10
AVAILABILITYZONE     |- m1.large      0008 / 0012    2     512     10
AVAILABILITYZONE     |- m1.xlarge     0008 / 0012    2    1024     20
AVAILABILITYZONE     |- c1.xlarge     0002 / 0006    4    2048     20
AVAILABILITYZONE     |- iwrcgblade11    certs[cc=false,nc=false]
                         @ Thu May 14 22:16:23 CEST 2009
AVAILABILITYZONE     |- iwrcgblade12    certs[cc=false,nc=false]
                         @ Thu May 14 22:16:23 CEST 2009
AVAILABILITYZONE     |- iwrcgblade13    certs[cc=false,nc=false]
                         @ Thu May 14 22:16:23 CEST 2009
AVAILABILITYZONE     |- iwrcgblade30    certs[cc=false,nc=false]
                         @ Thu May 14 22:16:23 CEST 2009
```

Register Images

Register a Filesystem-Image:

```
# ec2-bundle-image -i debian5.img
# ec2-upload-bundle -b image-debian5 -m /tmp/debian5.img.manifest.xml
# ec2-register image-debian5/debian5.img.manifest.xml
```

Register a Kernel-Image:

```
# ec2-bundle-image -i /boot/vmlinuz-2.6.26 --kernel true
# ec2-upload-bundle -b kernel26 -m /tmp/vmlinuz-2.6.26.manifest.xml
# ec2-register kernel26/vmlinuz-2.6.26.manifest.xml
```

Register a Ramdisk-Image:

```
# ec2-bundle-image -i /boot/initrd.img-2.6.26 --ramdisk true
# ec2-upload-bundle -b ramdisk26 -m /tmp/initrd.img-2.6.26.manifest.xml
# ec2-register ramdisk26/initrd.img-2.6.26.manifest.xml
```

Control Registered Images

- Registered images gets an unique identifier
  - Eucalyptus Machine Image: `emi-xxxxxxxx`
  - Eucalyptus Kernel Image: `eki-xxxxxxxx`
  - Eucalyptus Ramdisk Image: `eri-xxxxxxxx`
- Get information about registered images:

```
# ec2-describe-images
IMAGE  emi-1DE4116D   debian5/debian5.img.manifest.xml
                      admin   available    public   x86_64   machine
IMAGE  eki-791612FF   kernel26/vmlinuz-2.6.26.manifest.xml
                      admin   available    public   x86_64   kernel
IMAGE  eri-CFBE1450   ramdisk26/initrd.img-2.6.26.manifest.xml
                      admin   available    public   x86_64   ramdisk
```

Keypair creation and running instances

- **Create and register keypair:**

```
# ec2-add-keypair mykey > mykey.private
# chmod 0600 mykey.private
# ec2-describe-keypairs
KEYPAIR mykey
   33:da:6e:13:96:e6:f7:3b:b7:34:a6:28:ba:2f:64:ab:83:70:ef:70
```

- **Run instances:**

```
# ec2-run-instances emi-1DE4116D --kernel eki-791612FF
                                 --ramdisk eri-CFBE1450
                                 -k mykey -n 2 -t m1.small
```

  - `-k <keypair>`
  - `-n <#instances>`
  - `-t <instance_type>`
  - `-z <availability_zone>`

Control Instances and use them

Check instances:

```
# ec2-describe-instances
RESERVATION   r-3DDE07D9    admin           default
INSTANCE      i-4901084F    emi-1DE4116D    0.0.0.0    141.52.166.160
running       mykey         0               m1.small
2009-05-13T13:50:37+0000    eki-791612FF    eri-CFBE1450
RESERVATION   r-42FA0732    admin           default
INSTANCE      i-463B08BE    emi-1DE4116D    0.0.0.0    141.52.166.161
running       mykey         0               m1.small
2009-05-13T13:50:10+0000    eki-791612FF    eri-CFBE1450
```

Login to an instance via SSH:

```
# ssh -i mykey.private 141.52.166.160
```

Terminate instances:

```
# ec2-terminate-instances i-4901084F i-463B08BE
```

When the Installation fails...

- **Logging data from Eucalyptus**
  `/opt/eucalyptus/var/log/eucalyptus/`
  - **Cloud-Controller, Cluster-Controller and Node-Controller have their own Logfiles**

`cloud-debug.log, cloud-output.log`	Log data of Cloud Controller (CLC)
`cc.log`	Log data of Cluster Controller (CC)
`nc.log`	Log data of Node Controller (NC)
`httpd-cc_error_log`	STDERR/STDOUT of commands, started by the CC
`httpd-nc_error_log`	STDERR/STDOUT of commands, started by the NC

- **Logging data from Xen Hypervisor**
  `/var/log/xen/`
  - **Basic experience with Xen helps a lot**
- **Eucalyptus user forum**
  `http://forum.eucalyptus.com/forum/`
- **Coffee, time and strong nerves**

Handy and popular tools for Eucalyptus

- The Amazon Web-Services (EC2, S3, EBS, ...) are very popular
- Lots of handy tools supporting the AWS exist
- Because of Eucalyptus' interface compatibility with Amazon EC2, S3 and EBS many of these tools can be utilized with Eucalyptus
  - Why not all? Because the API-Release differs!
- Handy tools
  - S3 Curl (command line tool)
    - `http://developer.amazonwebservices.com/connect/entry.jspa?externalID=128`
  - s3cmd (command line tool)
    - `http://s3tools.org/s3cmd`
  - s3fs (S3-Bucket can be mounted as local filesystem with FUSE)
    - `http://code.google.com/p/s3fs/`
  - ElasticFox (Firefox-Plugin)
    - `https://code.launchpad.net/~soren/elasticfox/elasticfox.eucalyptus`

ElasticFox

- Support to start, monitor and terminate instances in a user friendly GUI
- Compatibility to Eucalyptus need to be improved
- Release 55 works quite well with Eucalyptus 1.5.1
- Eucalyptus 1.6 (September 2009) shall help a lot
  - Compatibility to Amazon AWS specification 1/1/2009

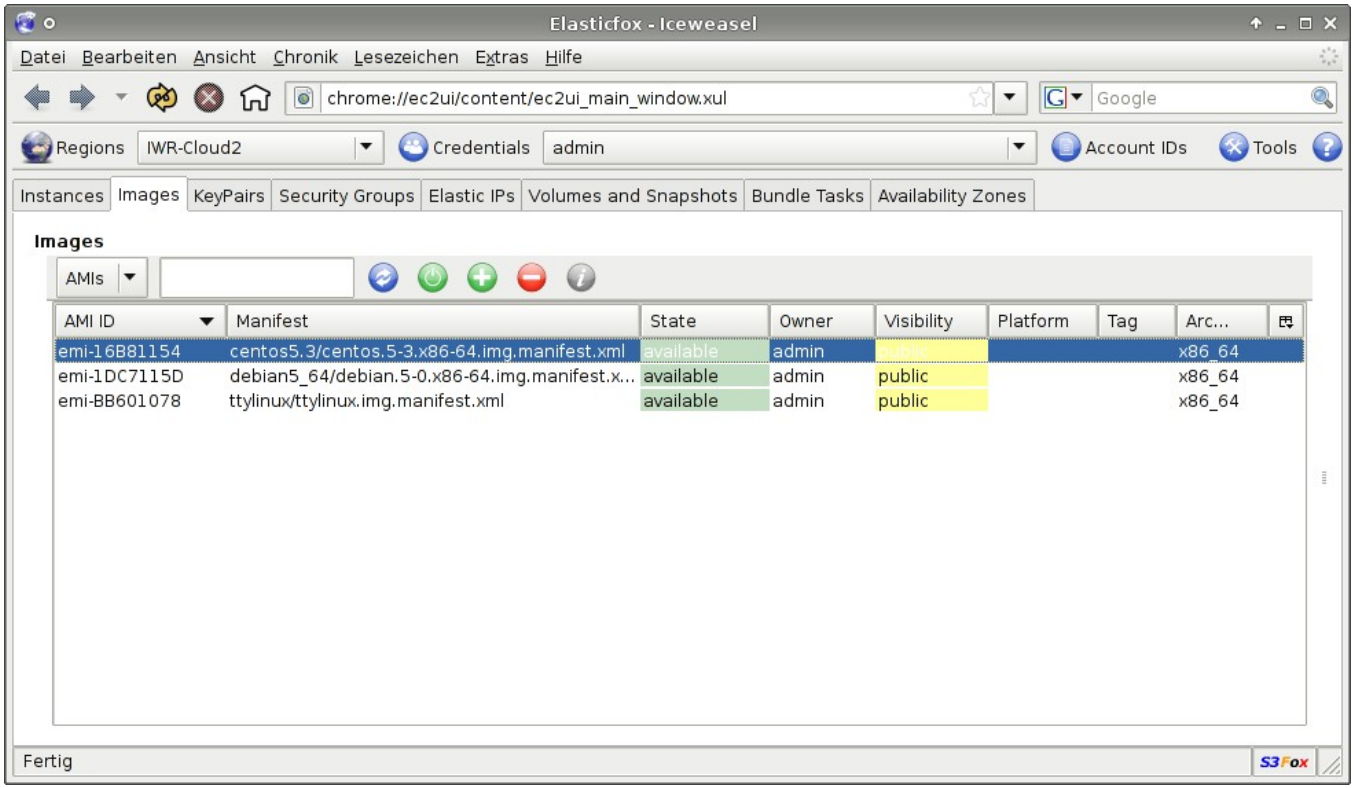

Storage Performance S3 vs. Eucalyptus

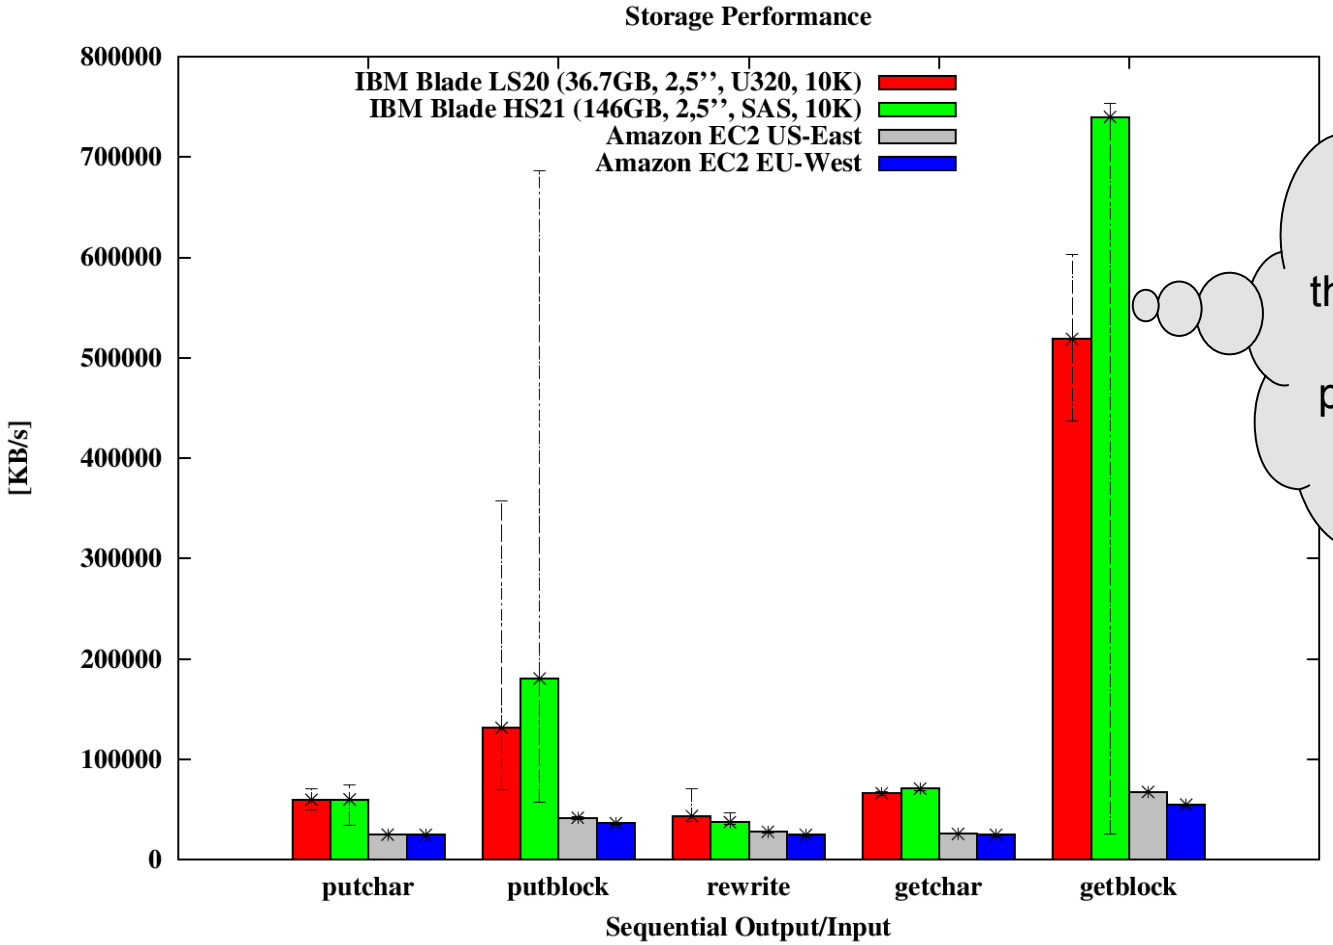

WOW!?
the blockwise read/write performance is fantastic

- **Sequential Output**
  - Per-Character: file is written using putc()
  - Blockwise: file is written using `write()`
  - Rewrite: `read()` and `write()`
- **Sequential Input**
  - Per-Character: file is read using `getc()`
  - Blockwise: file is read usind `read()`

Realistic values…

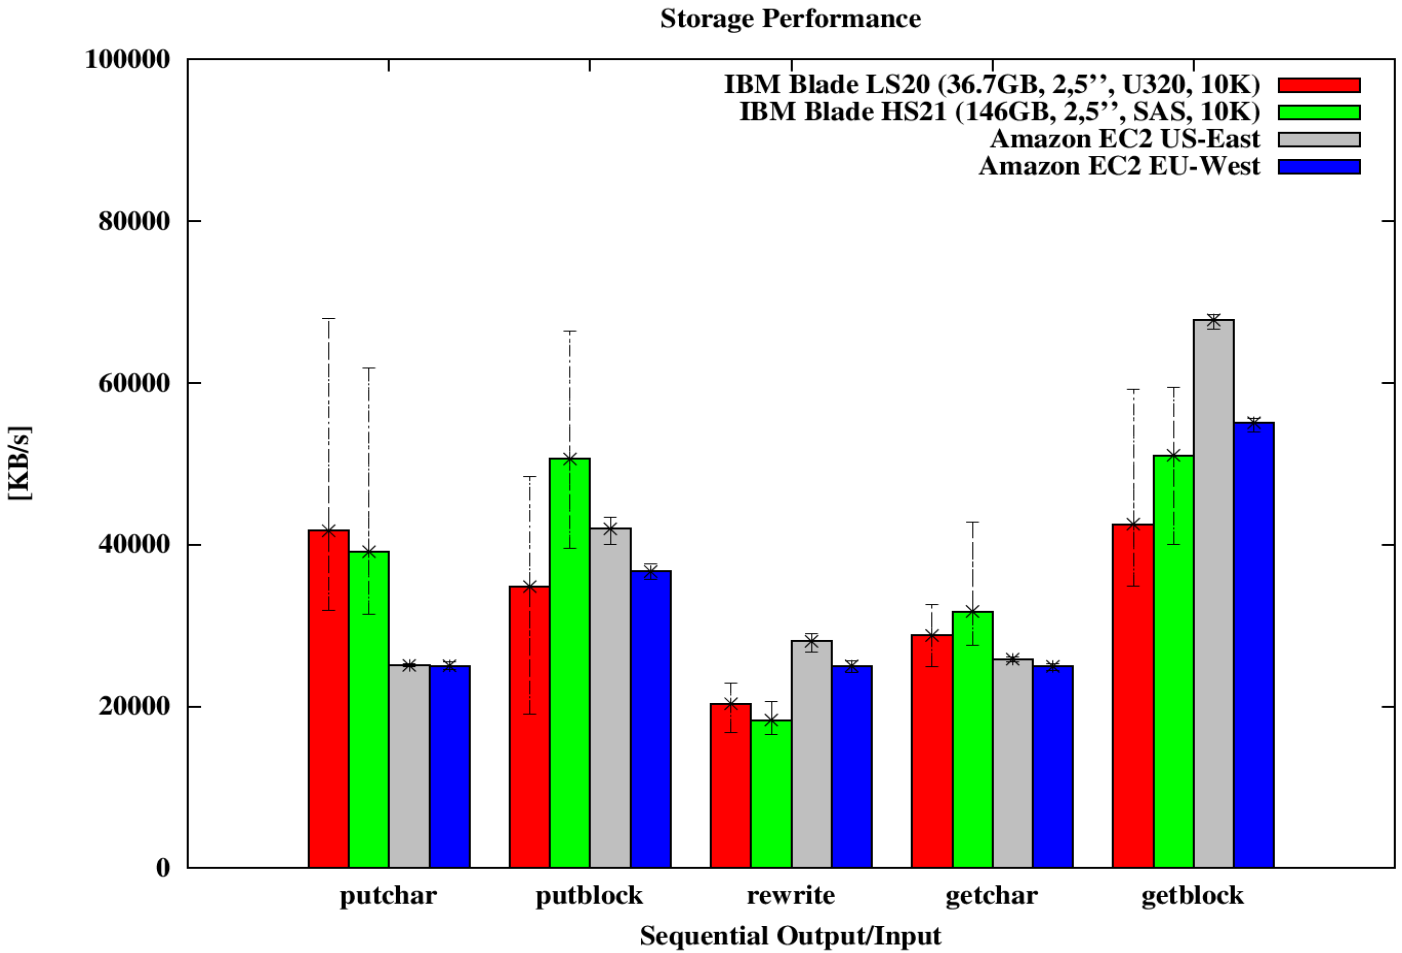

- The RAM of the Eucalyptus Node Controller was reduced to overcome the memory caching of the Linux kernel
- The write performance of Eucalyptus is better in this scenario
- Eucalyptus' storage performance depends on the available storage sub-system

Random Seeks and File Creation/Deletion

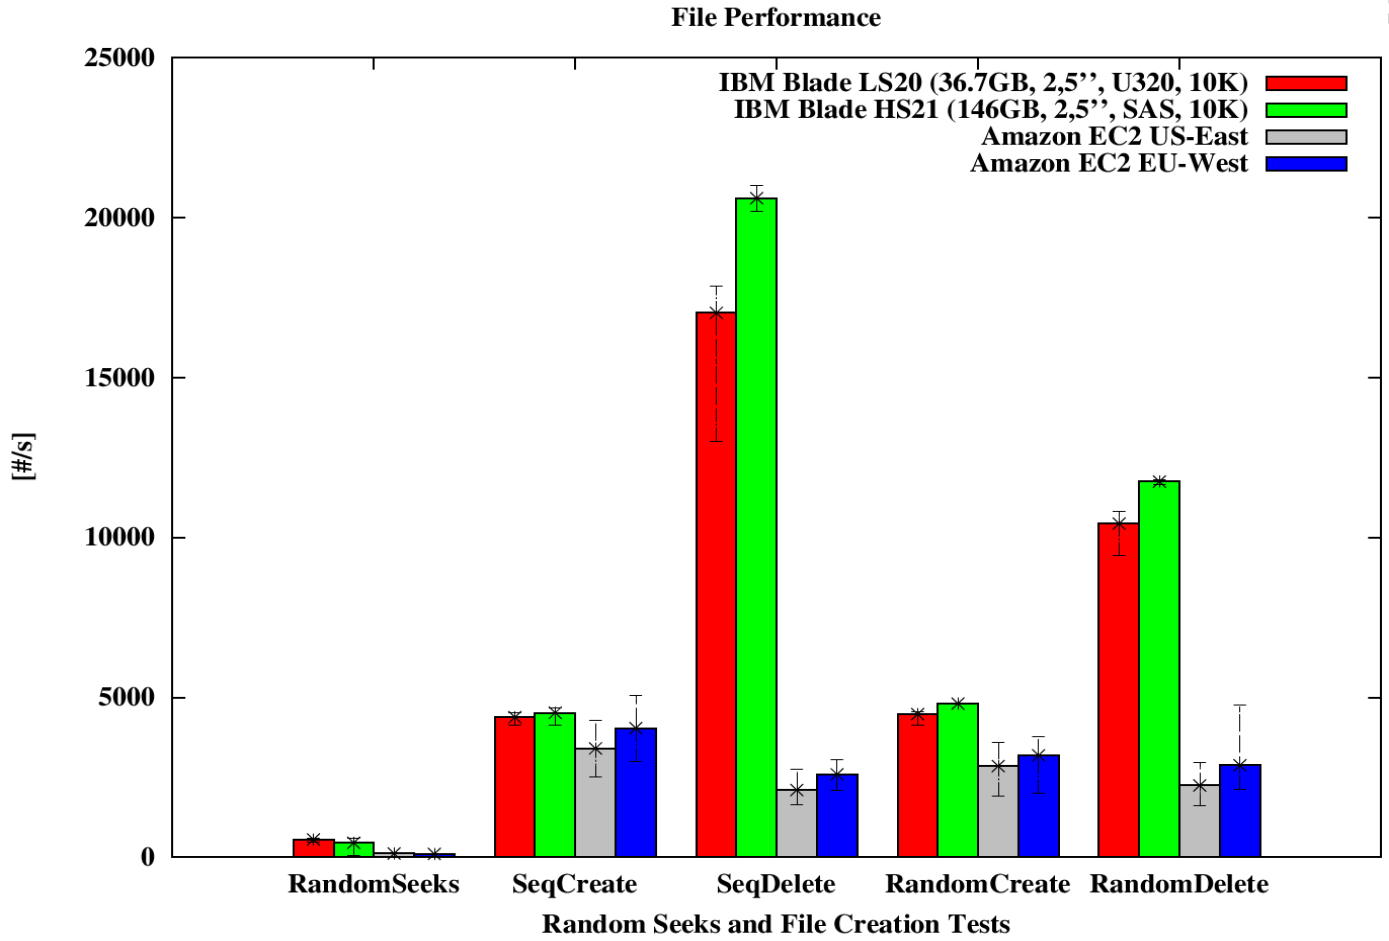

- Random seeks and file creation/deletion with Eucalyptus is faster compared to Amazon S3
- Because of the close distance between the EC2 and S3 components?
- Performance capability and workload of Amazon S3 is unknown

CPU Performance

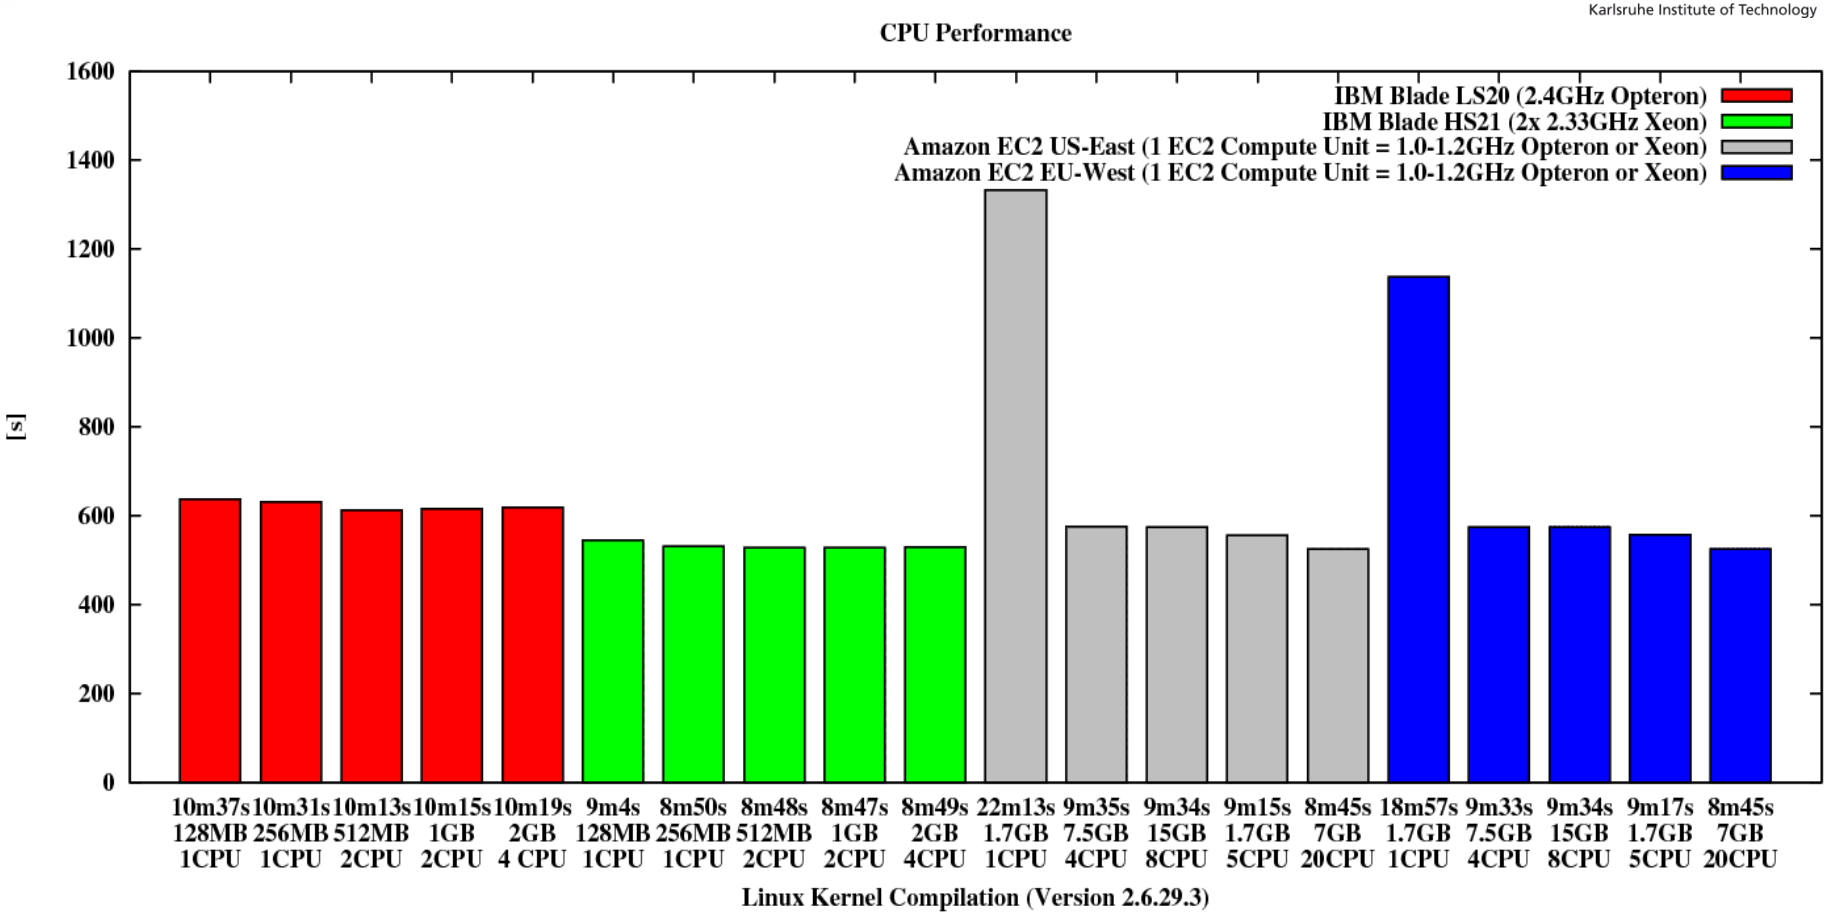

- The performance of the Eucalyptus infrastructure (instance types) is comparable with Amazon EC2
- Interesting: More RAM and CPU power is not leading to a massive performance boost necessarily

Network Transfer Rate

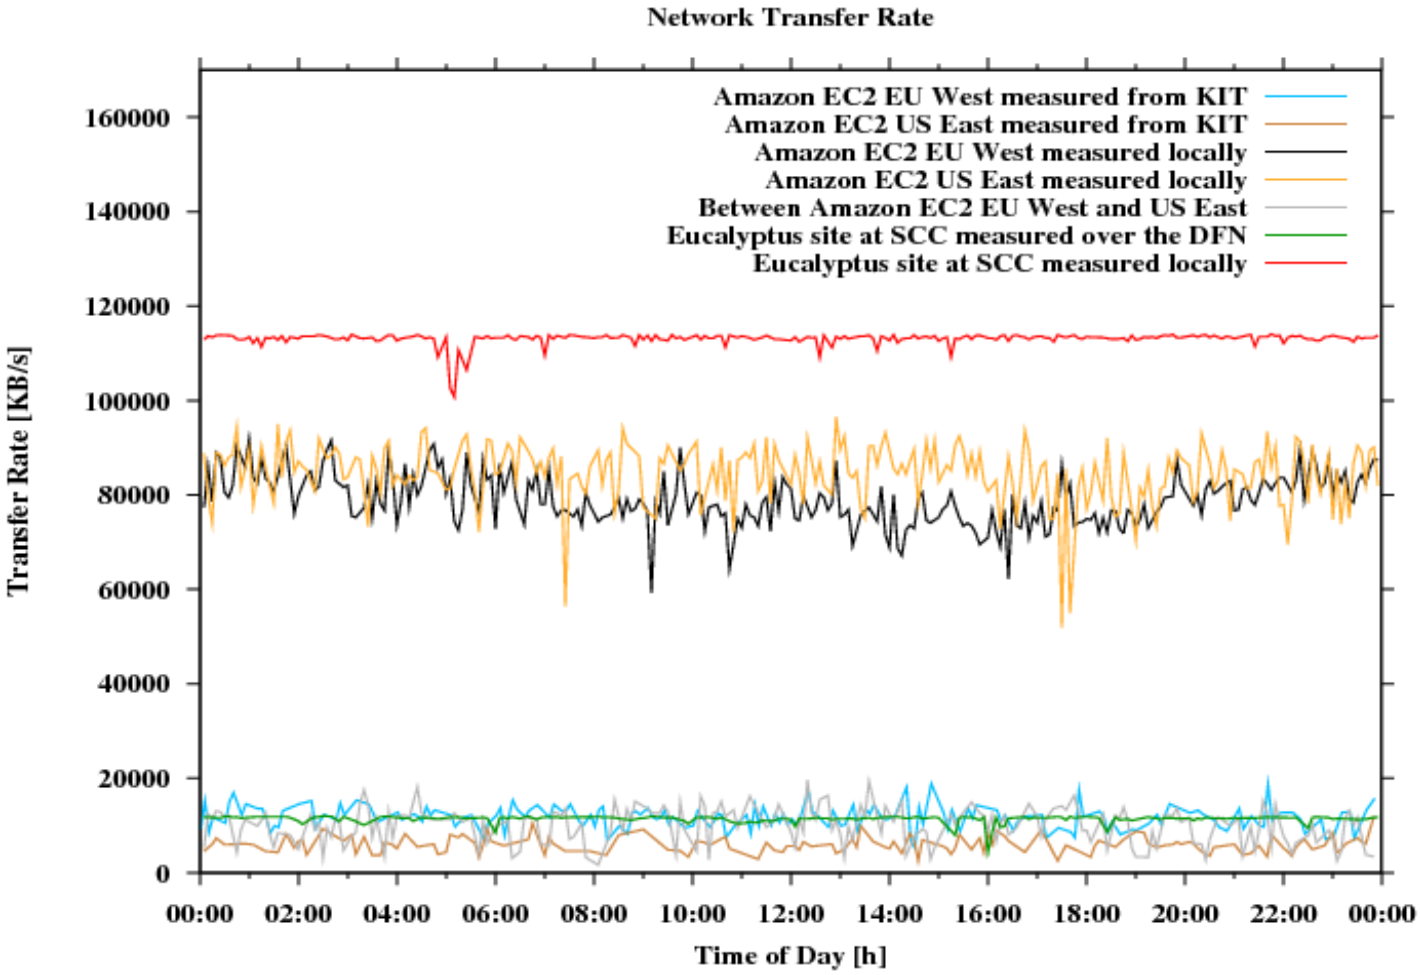

- **Network Transfer Rate from/to Eucalyptus at FZK and Amazon EC2 (also measured inside)**
- **This graph is confusing ...**

Network Transfer Rate - More in Detail (1)

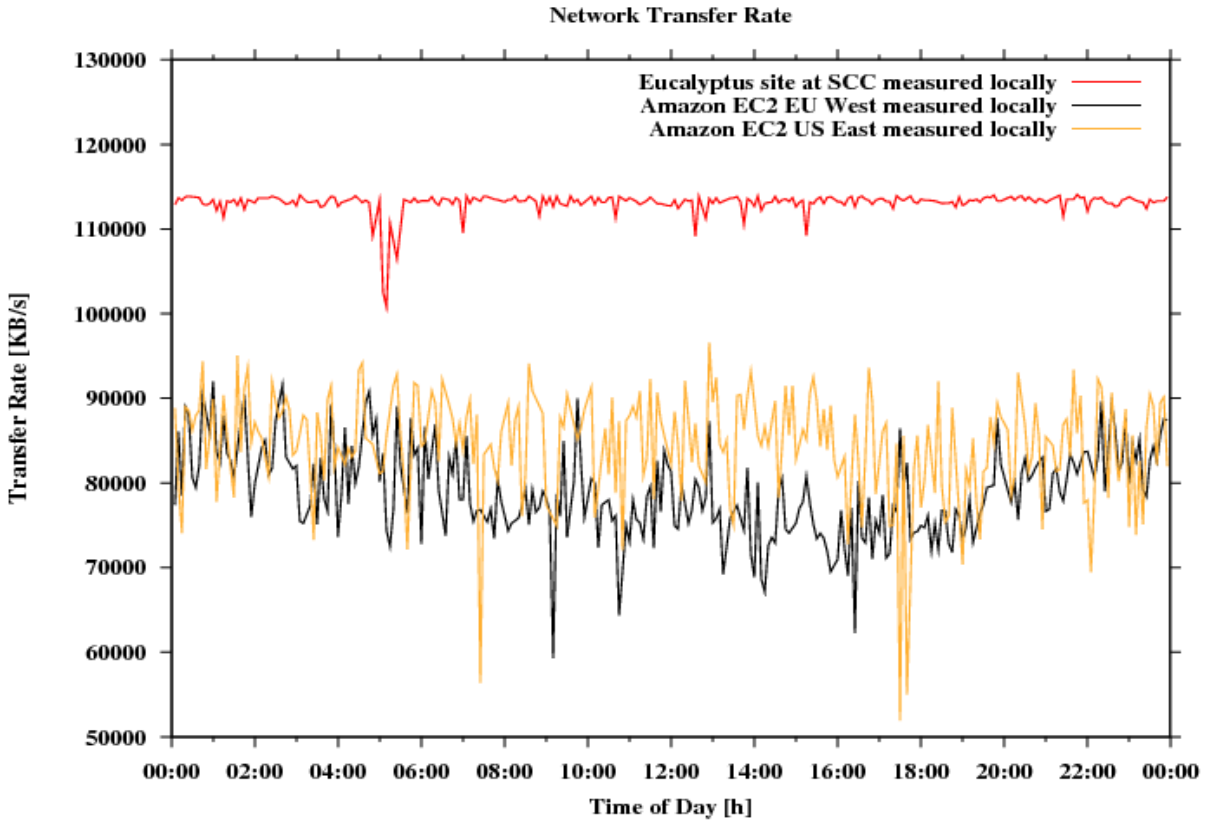

- Network Transfer Rate within the Eucalyptus site is more constant compared to Amazon EC2
- The reason for the strong Network Transfer Rate between the Eucalyptus Nodes is the 1000 Mbit/s Ethernet
- We assume Amazon uses 1000 Mbit/s Ethernet in their EC2-Sites too but there is much more workload

Network Transfer Rate - More in Detail (2)

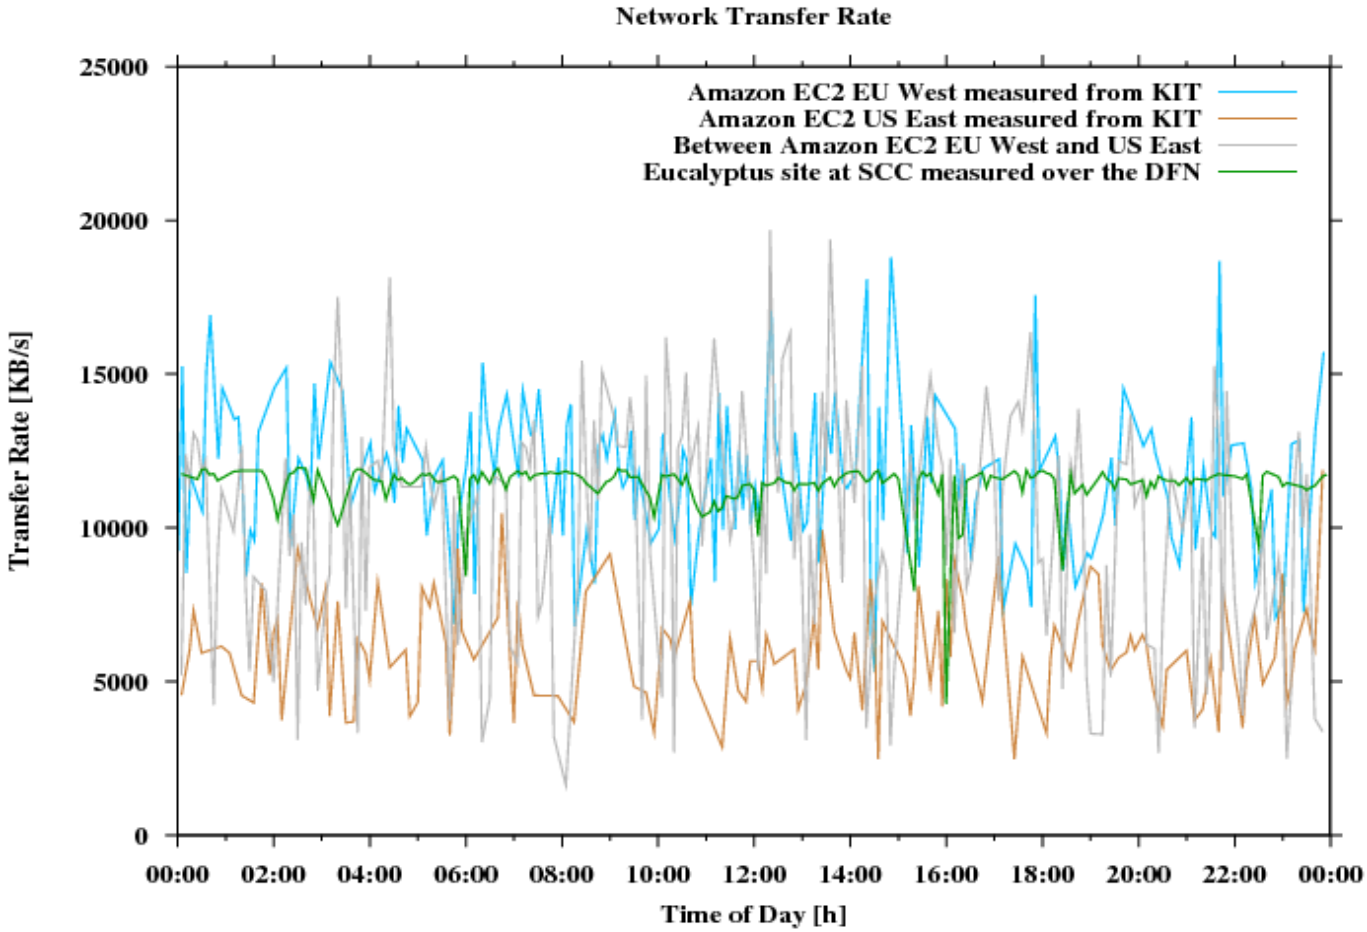

- The Network Transfer Rate from inside Europe to Amazon EC2 EU is much better compared to Amazon EC2 US. This is not surprising
- The Network Transfer Rate to Eucalyptus over DFN is more constant compared to Amazon EC2

AppScale

http://appscale.cs.ucsb.edu

- Open Source reimplementation of the Google AppEngine
  - AppEngine allows to run web applications written in Python (and JAVA) in the Google infrastructure
- Developed at UC Santa Barbara
- Implements Platform as a Service (PaaS)
- AppScale executes automatically and transparently in Eucalyptus
- With AppScale a PaaS Cloud infrastructure can be build up that allows to deploy, test, debug, measure, and monitor Google AppEngine applications inside a Private Cloud

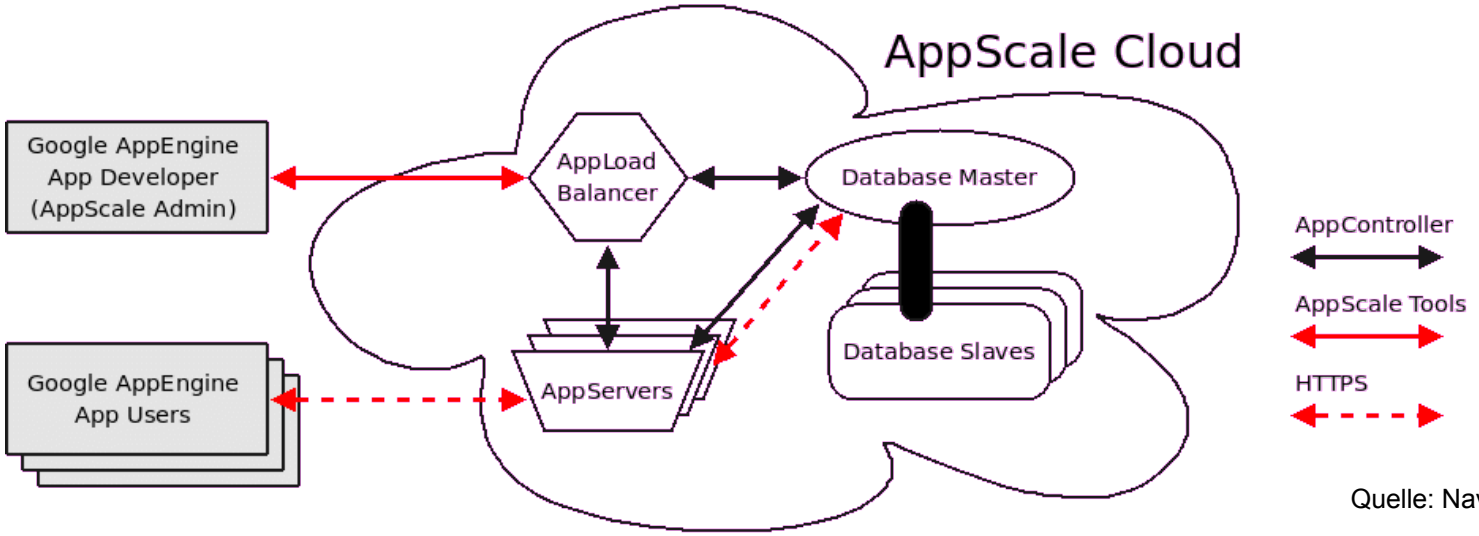

Quelle: Navraj Chohan

OpenCirrus™ In the Press

Karlsruhe Institute of Technology

InformationWeek
BUSINESS INNOVATION POWERED BY TECHNOLOGY

HP, Intel, Yahoo Join Government, Academia In Cloud Computing Research

Each of the founding members will host a cloud-computing infrastructure largely based on HP computers and Intel processors in six data centers.

By Antone Gonsalves, InformationWeek
July 29, 2008
URL: http://www.informationweek.com/story/showArticle.jhtml?articleID=209800449

Hewlett-Packard, Intel, and Yahoo on Tuesday said they have joined government and academia in launching a global, multi-data center test bed for experimentation and research in cloud computing, which many experts believe will be the dominant IT delivery model of the future.

The initiative aims at building a computing network comprised of six data centers spanning three continents. The idea is to have a large-scale platform for testing all technology -- hardware and software -- related to delivering application services over the Internet.

"This is a global collaboration that spans the industry, spans academia and government," Prith Banerjee, senior VP for research at HP, told reporters during a teleconference held by the three founding companies.

The other founders of the effort include the Infocomm Development Authority of Singapore, the University of Illinois at Urbana-Champaign, and the Karlsruhe Institute of Technology in Germany. The partnership with the University of Illinois also includes the National Science Foundation.

Each of the founding members will host a cloud-computing infrastructure largely based on HP computers and Intel processors. The infrastructure will include from 1,000 to 4,000 processor cores capable of supporting data-intensive research. The six facilities are up and running today in "bits and pieces" and are expected to be fully operational this year and accessible to researchers worldwide through a selection process.

OpenCirrus™ Cloud Computing Research Testbed

- An open, internet-scale global testbed for cloud computing research
  - Data center management & cloud services
  - Systems level research
  - Application level research
- Structure: a loose federation
  - Sponsors: HP Labs, Intel Research, Yahoo!
  - Partners: University of Illinois at Urbana-Champaign (UIUC), Singapore Infocomm Development Authority (IDA), KIT
- Great opportunity for cloud R&D
- `http://opencirrus.org`

Where are the OpenCirrus™ sites?

- Six sites initially:
  - Sites distributed world-wide: HP Research, Yahoo!, University of Illinois (UIUC), Intel Research Pittsburgh, KIT, Singapore IDA
  - 1000 - 4000 processor cores per site
- KIT-Site available in Summer 2009
  - 3300 Nehalem cores, 10TB memory, 192TB hard disk storage

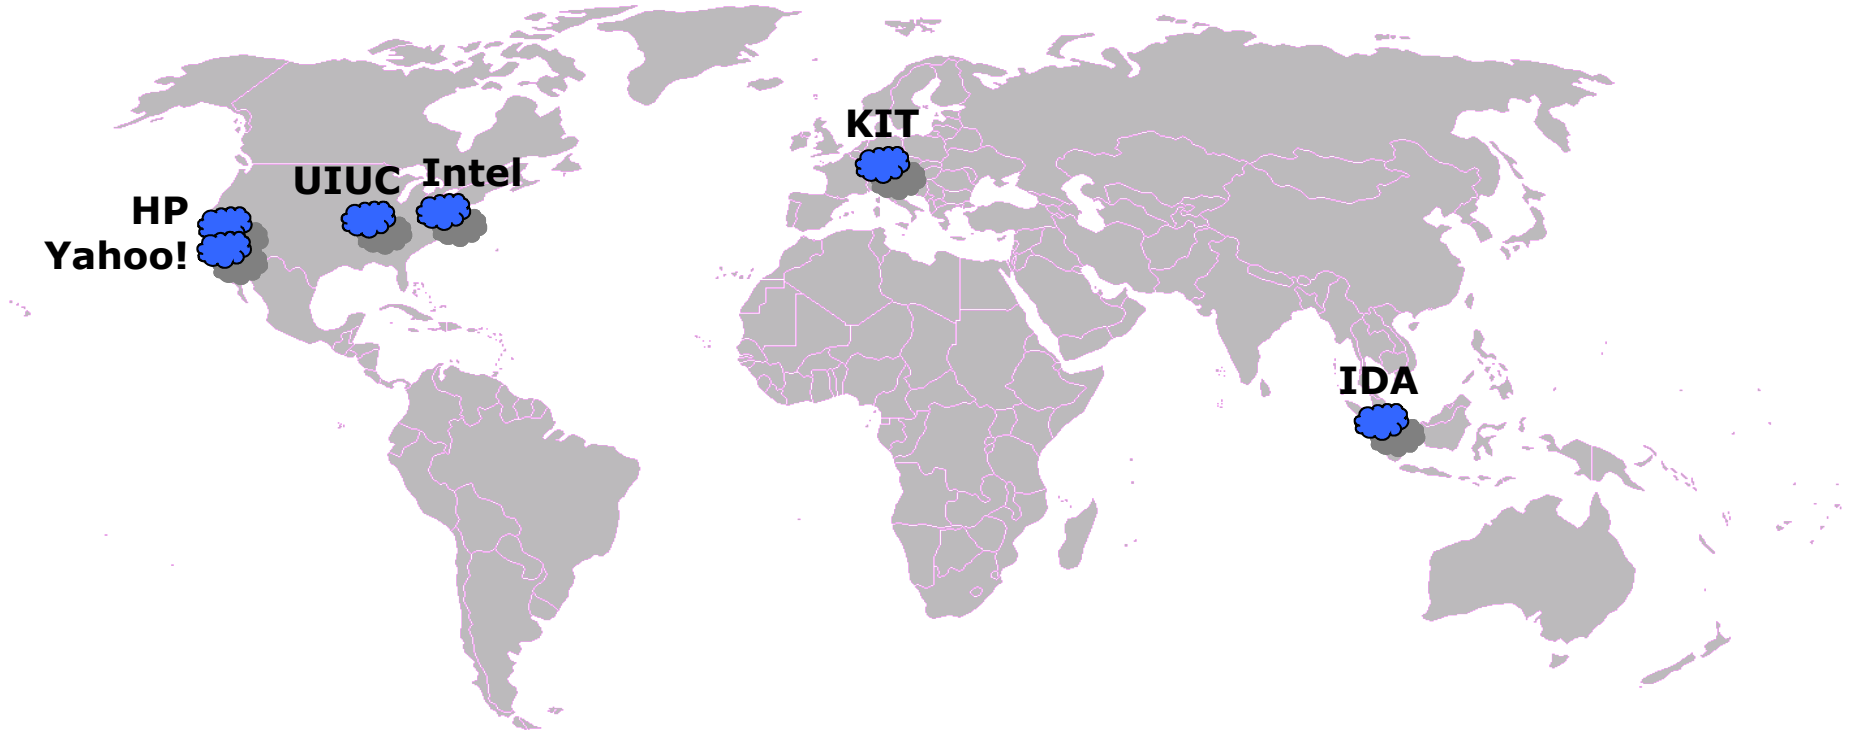

Summary

- Cloud computing is the next big thing
- Promising approach to solve some of the major challenges of IT
- Flexible and elastic resource provisioning
- Economy of scale makes it attractive
- Move from manufacture towards industrialization of IT
- Eucalyptus and AppScale enable IaaS und PaaS als Open Source solutions with Linux
- But: In the Cloud, there is still much to do

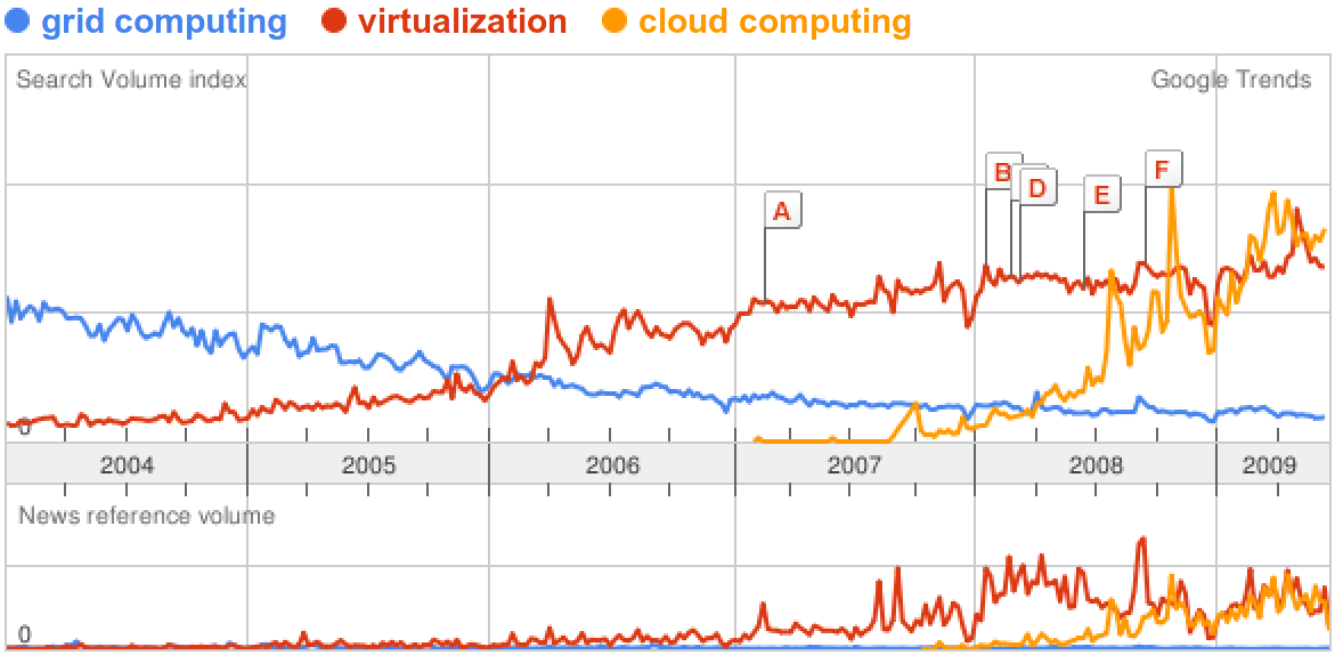

Coming soon - this autumn

- Available in October 2009
- First Cloud Computing book in German language
- Covers the latest topics in Cloud Computing
- Only € 14,95

Thank you for your attention