Karlsruhe Institute of Technology

Performance measurement of a private Cloud in the OpenCirrus™ Testbed

4th Workshop on Virtualization in High-Performance Cloud Computing (VHPC '09)

Euro-Par 2009 – Delft – August 25th 2009

Christian Baun

Forschungszentrum Karlsruhe
in der Helmholtz-Gemeinschaft

Universität Karlsruhe (TH)
Forschungsuniversität • gegründet 1825

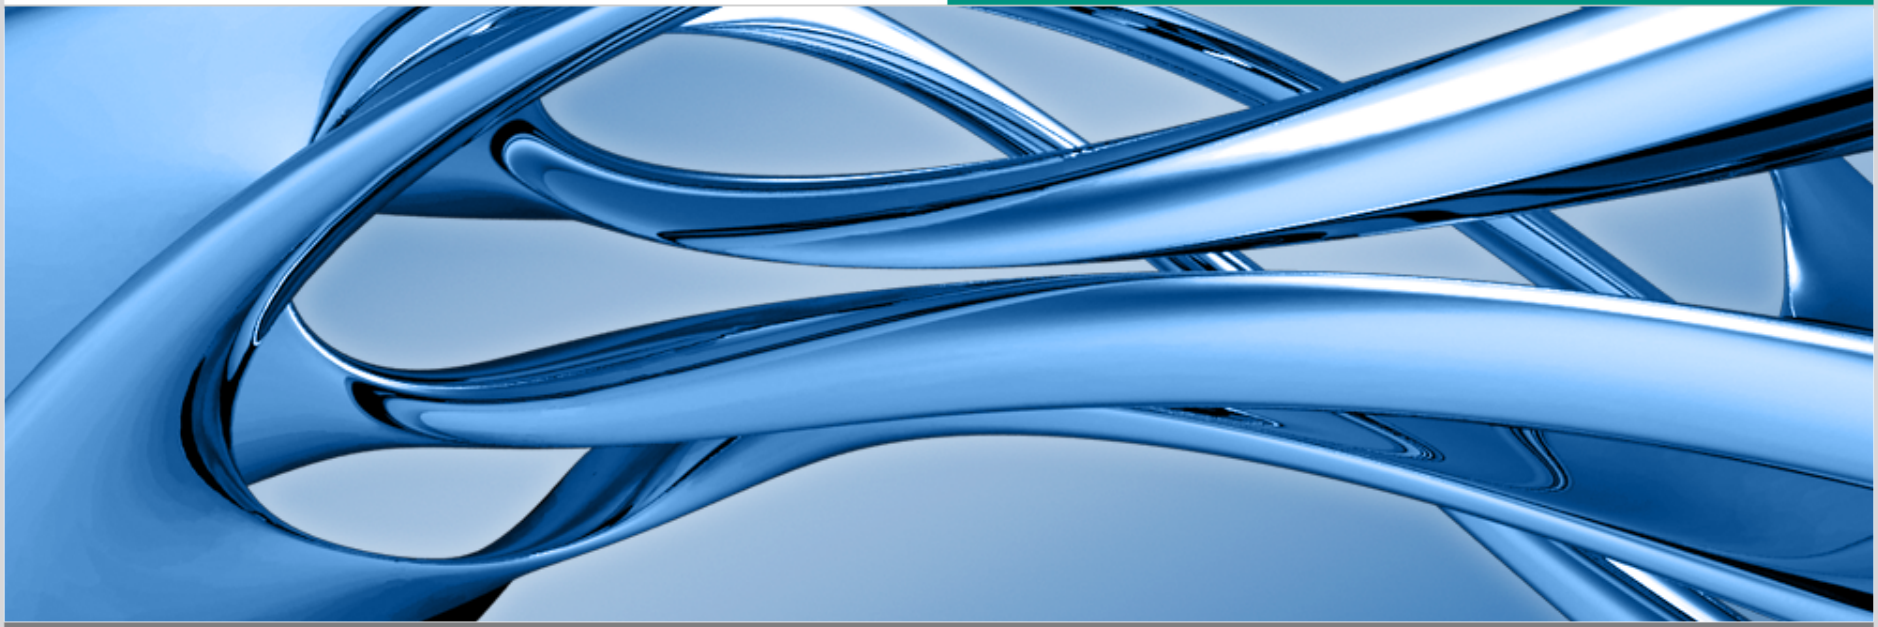

Cloud Computing - big topic/hype

- 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

Web 2.0: Article

Cloud Hype at Height: Gartner

It's supposed to be 2–5 years away from mainstream adoption

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.

Why do we want/need Cloud Computing?

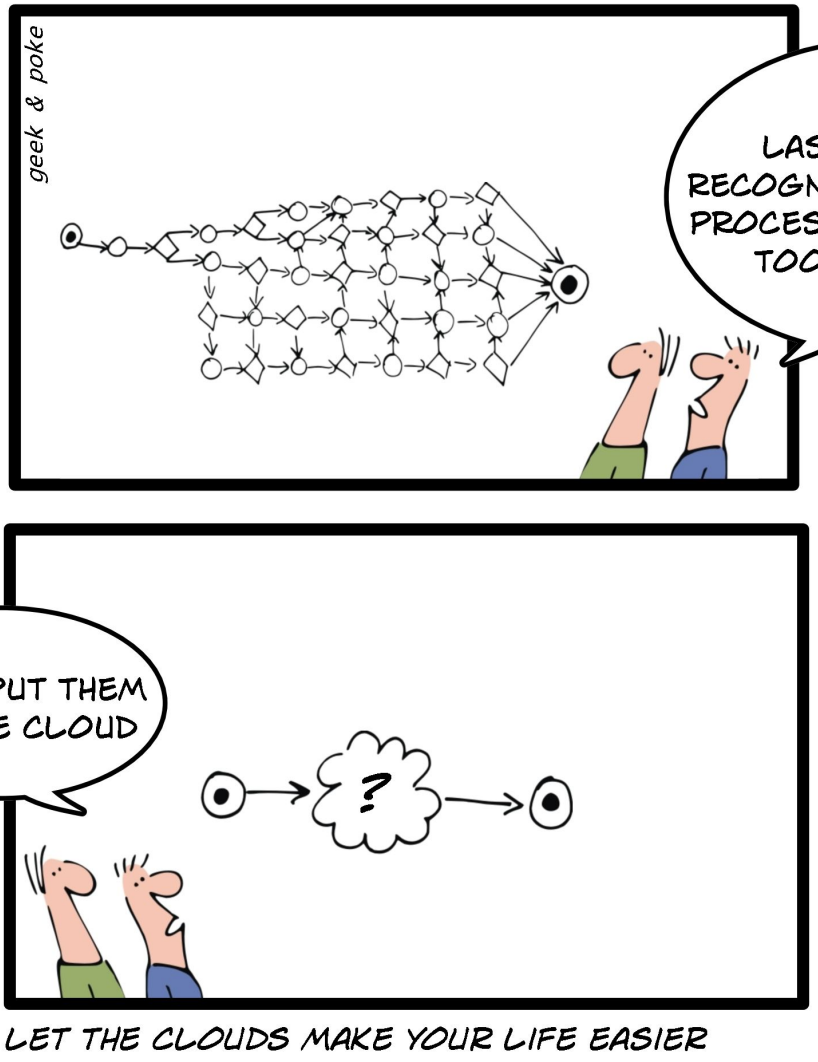

- Cloud Computing shall:
  - revolutionize IT
  - reduce (eliminate) complexity
  - getting IT more flexible (elastic)
  - reduce cost
  - provide easy resources access
  - fulfill users demands
  - emancipate the users
  - ...
- Remember: GRID !?!?!?

Definitions of Cloud-Computing

- **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)

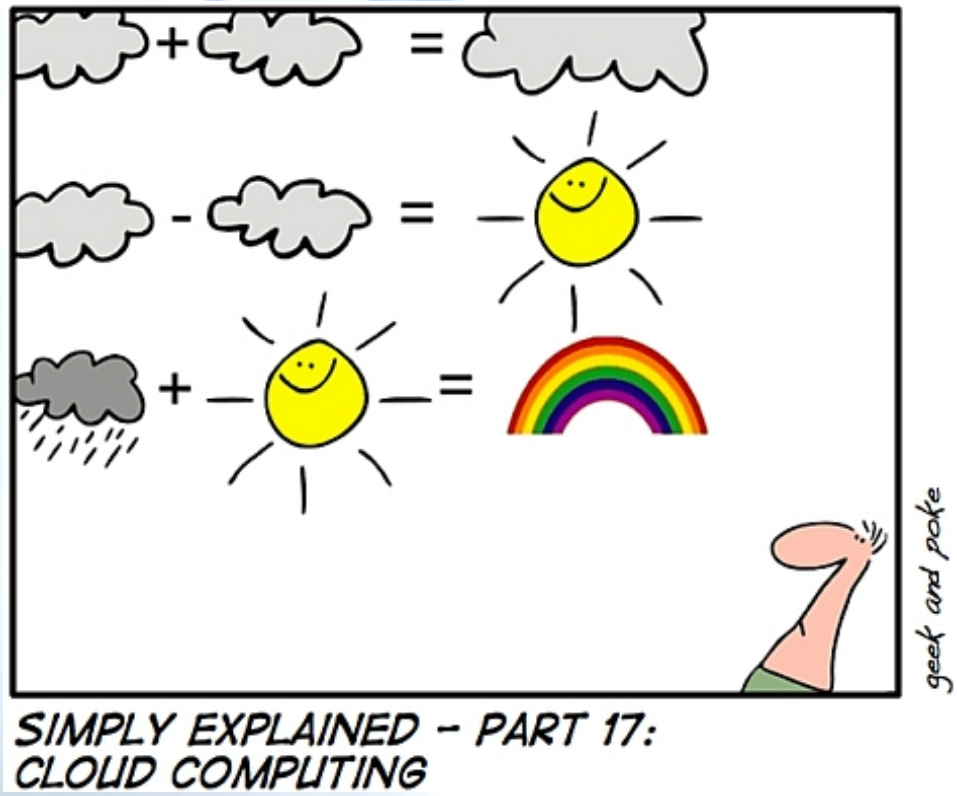

Three technical Types of Cloud Services

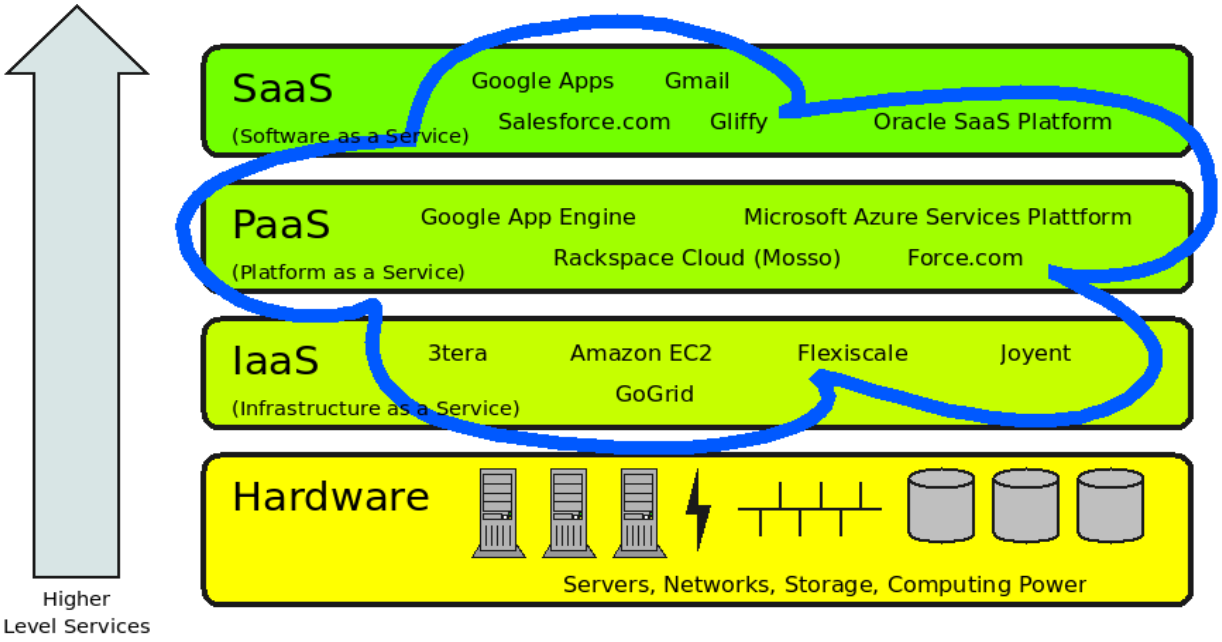

- **SaaS**
  - **Provides enterprise quality software that is run by a service provider**
  - **The user has no need to worry about installation, administration or updates**
- **PaaS**
  - **Appears as one single large computer and makes it simple to scale from a single server to many**
  - **No need 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 Infrastructures

- 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

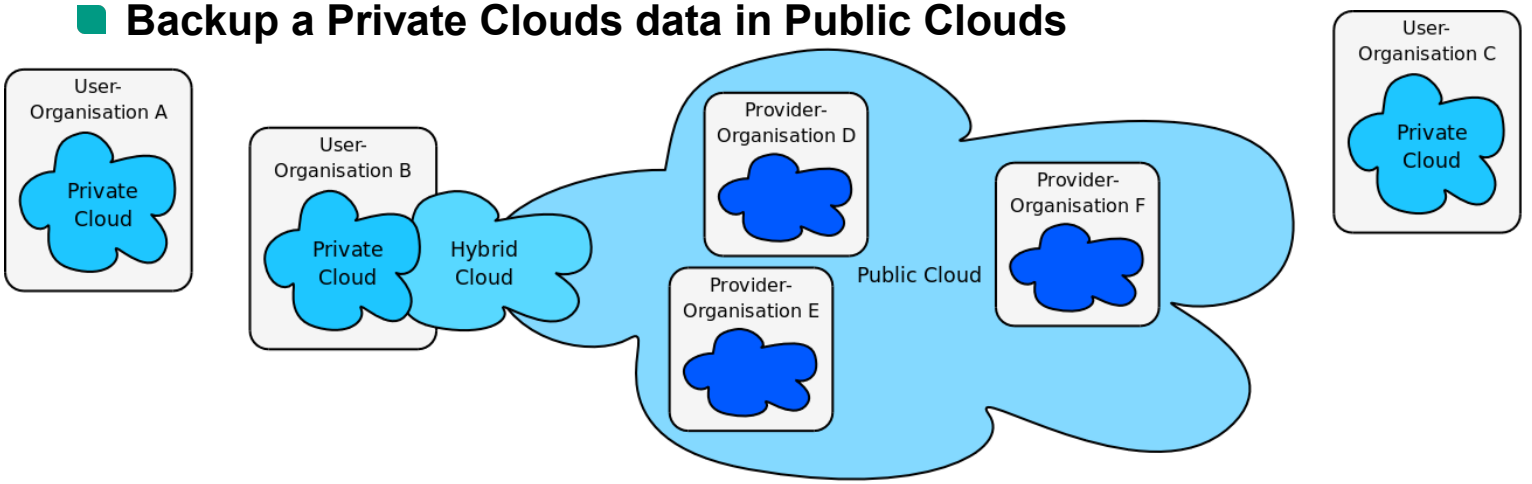

OpenCirrus™ In the Press

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!, UIUC, Intel Research Pittsburgh, KIT, Singapore IDA
  - 1000 - 4000 processor cores per site
- KIT-Site available in Autumn 2009
  - 3300 Nehalem cores, 10TB memory, 192TB hard disk storage

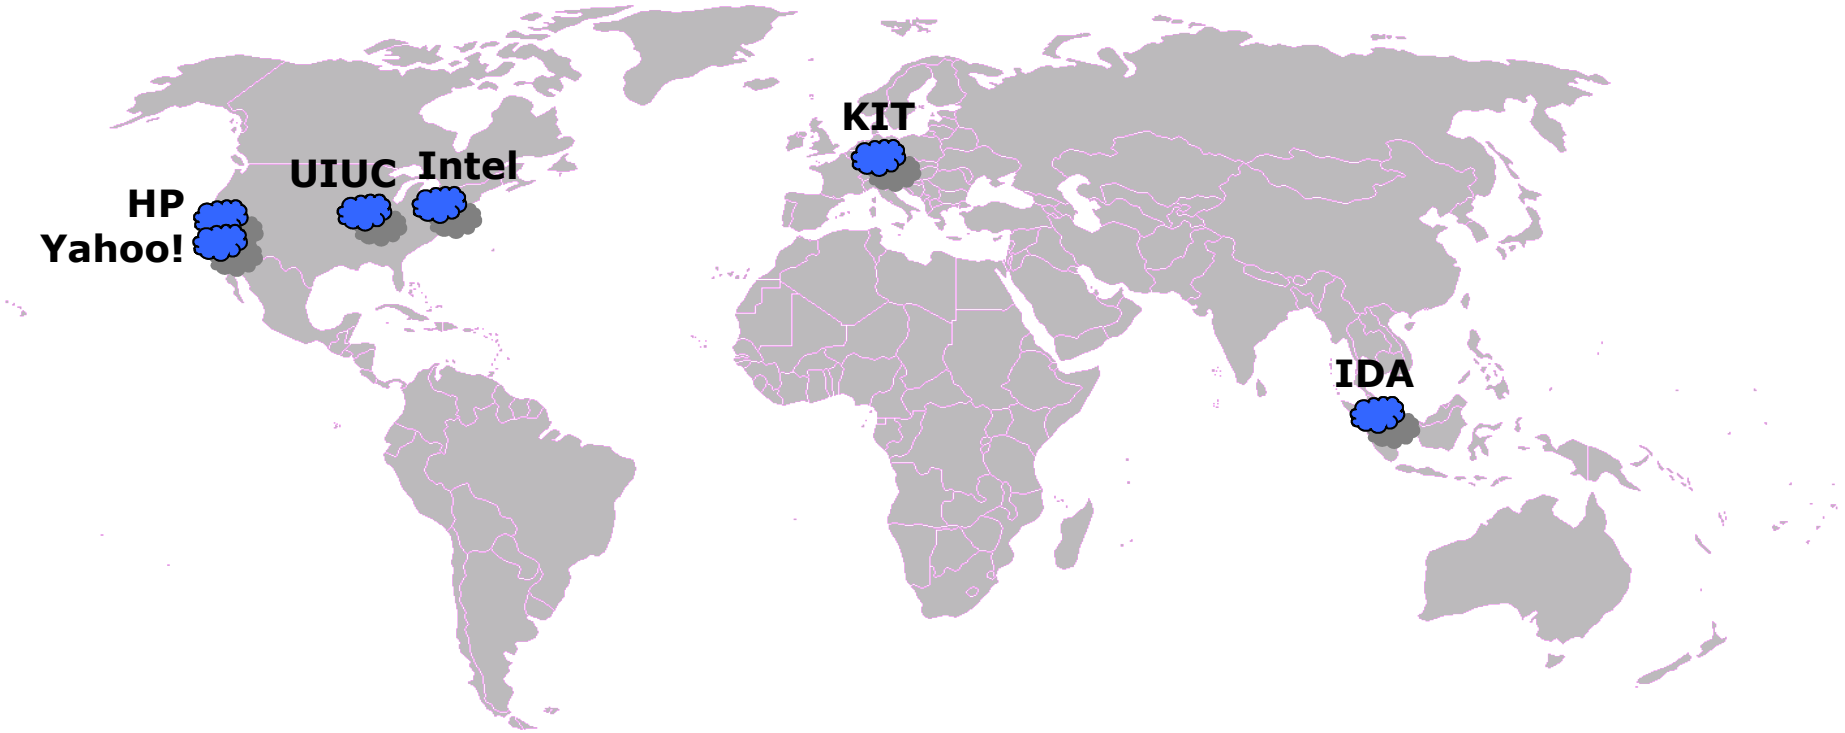

OpenCirrus™ - Physical Resource Sets (PRS)

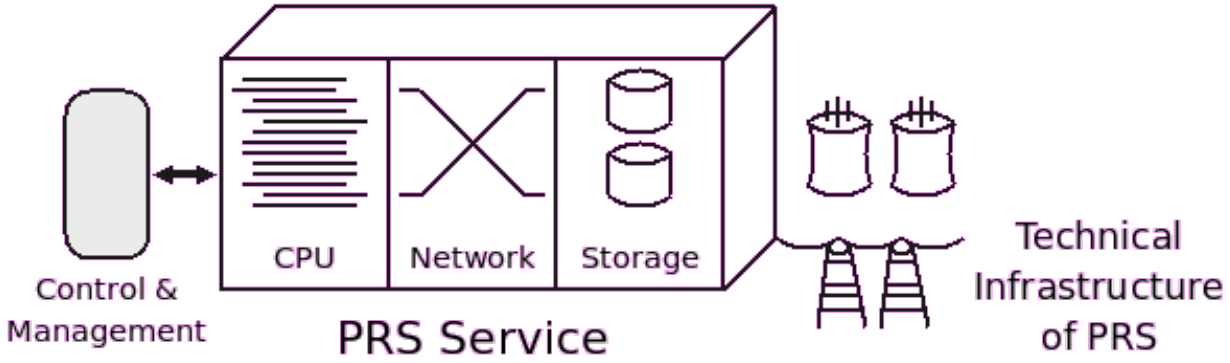

- PRS service goals
  - Provide mini-datacenters to the researchers
  - Isolate experiments from each other
- PRS service approach
  - Allocate sets of physical co-located nodes, isolated inside VLANs using existing software
    - Utah Emulab - Network Emulation Testbed
    - HP Opsware - Server provisioning, configuration and management
  - Start simple, add features as we go
  - Cloud Infrastructure Services (e.g. EC2/Eucalyptus, S3 Storage, NFS Storage) are implemented via PRS and partly provided as virtual instances
    - Virtual Resource Sets (VRS)
- Hardware as a Service (HaaS)

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://open.eucalyptus.com

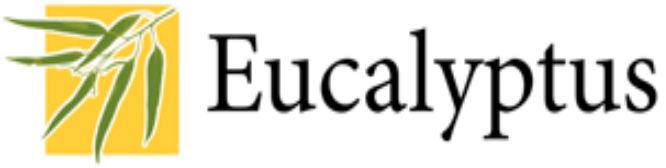

- EUCALYPTUS - Elastic Utility Computing Architecture for Linking Your Programs To Useful Systems
- Open Source software infrastructure for implementing Cloud Computing on clusters from UC Santa Barbara
- Developed at UC Santa Barbara
- Implements Infrastructure as a Service (IaaS)
- Gives the user the 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 popular tools, known to work with Amazon EC2, S3 and EBS e.g. S3 Curl, ElasticFox, s3cmd, ...
- Eucalyptus is an important step to establish an open Cloud Computing infrastructure standard

Eucalyptus - Components

http://open.eucalyptus.com

Eucalyptus

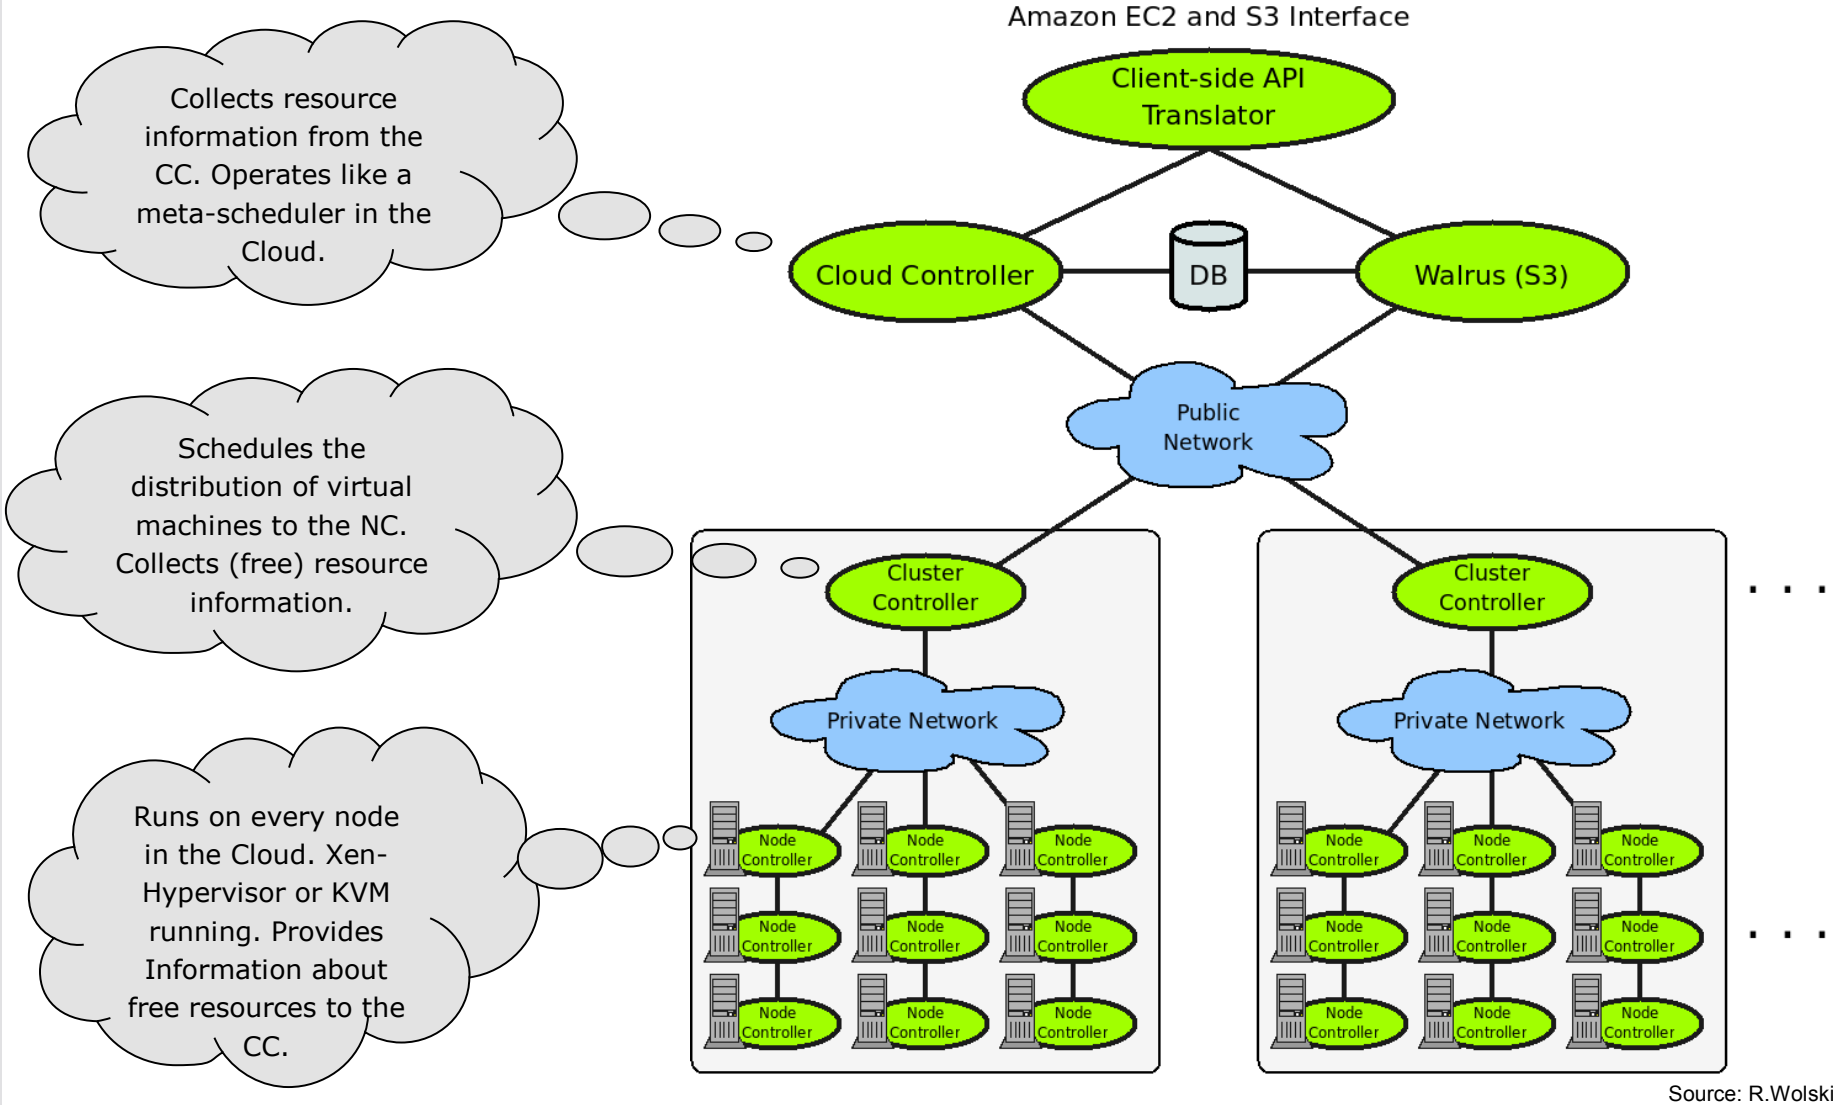

Source: R.Wolski

Eucalyptus at KIT

- Private 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
  - Runs stable at KIT since February 2009

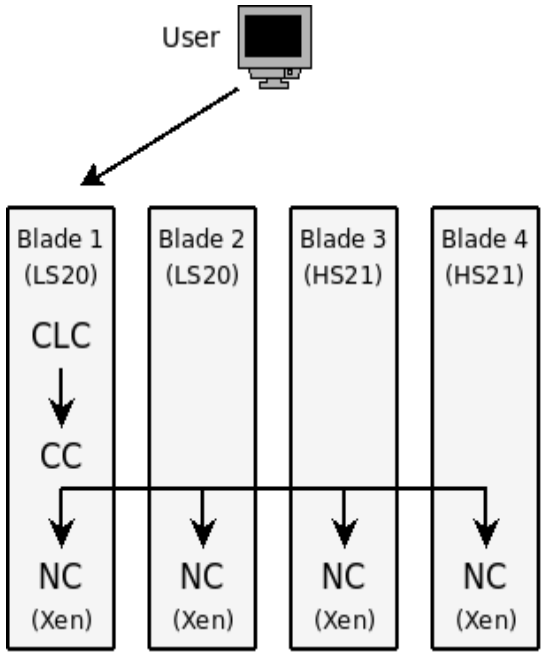

- Private Cloud Installation II (Eucalyptus 1.5.2)
  - Under construction
  - 5x HP Blade ProLiant BL2x220c
  - 2 Server per blade:
    - 2x Intel Quad-Code Xeon (2,33GHz)
    - 16GB RAM

- Question: What is the performance compared to Amazon EC2/S3?

Storage Performance S3 vs. Eucalyptus

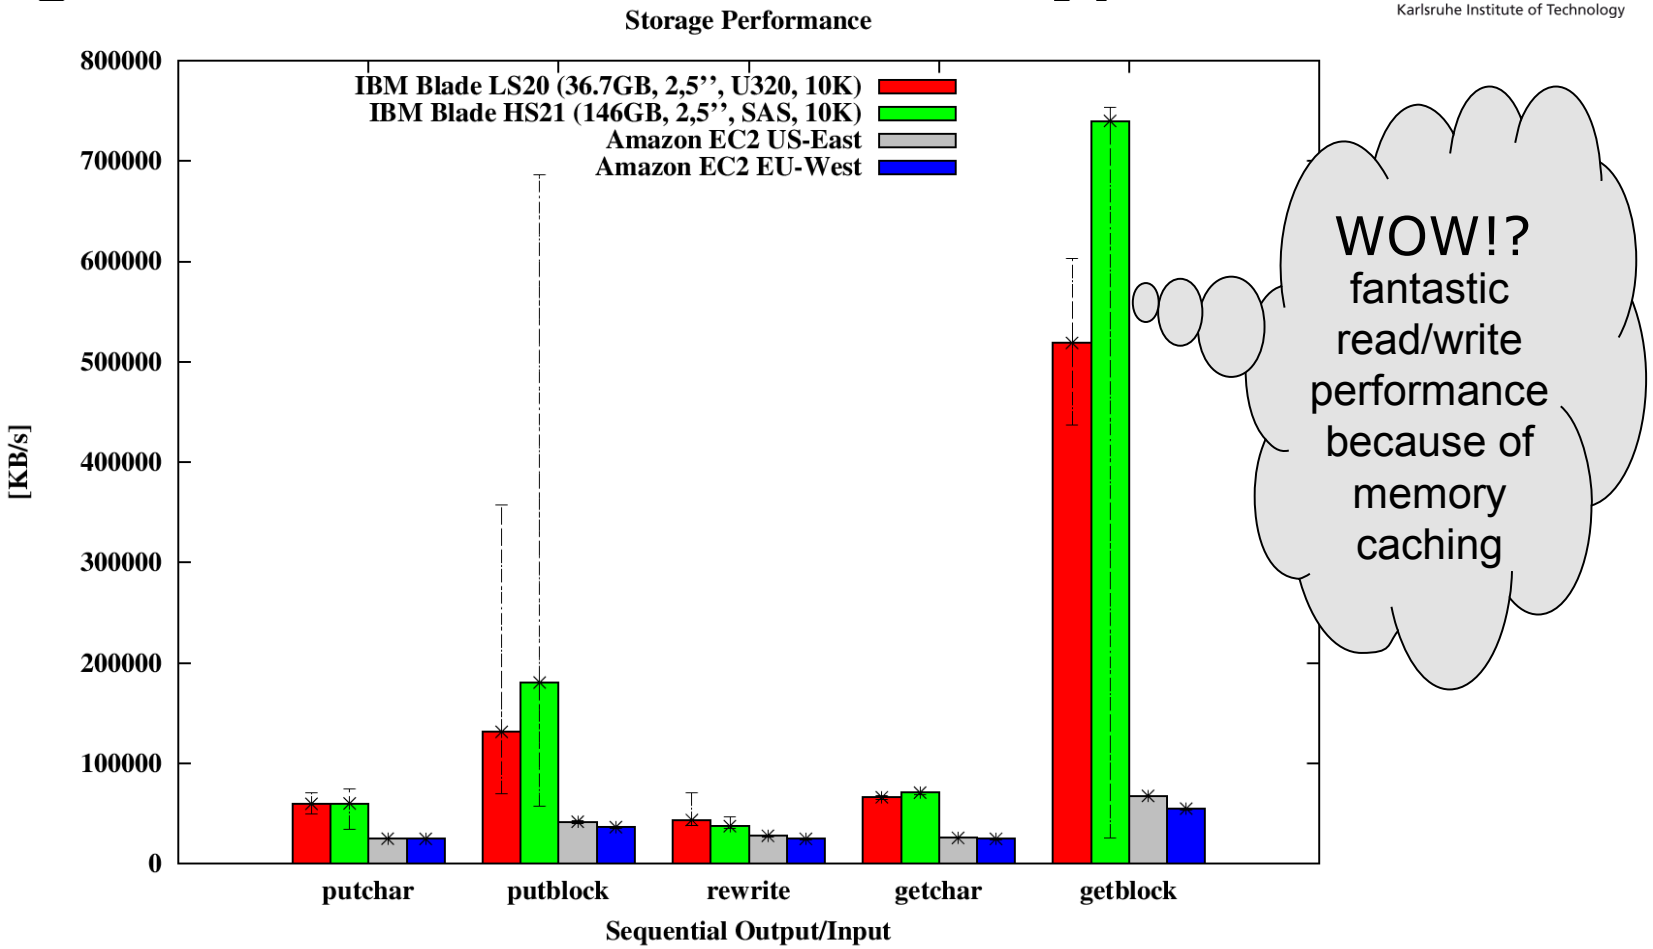

- **Sequential Output**
  - Per-Character: file is written using putc()
  - Blockwise: file is written using `write()`
  - Rewrite: `read()` and `write()`
- **Sequential Input**
  - Per-Character: file is read using `getc()`
  - Blockwise: file is read usind `read()`

Realistic values…

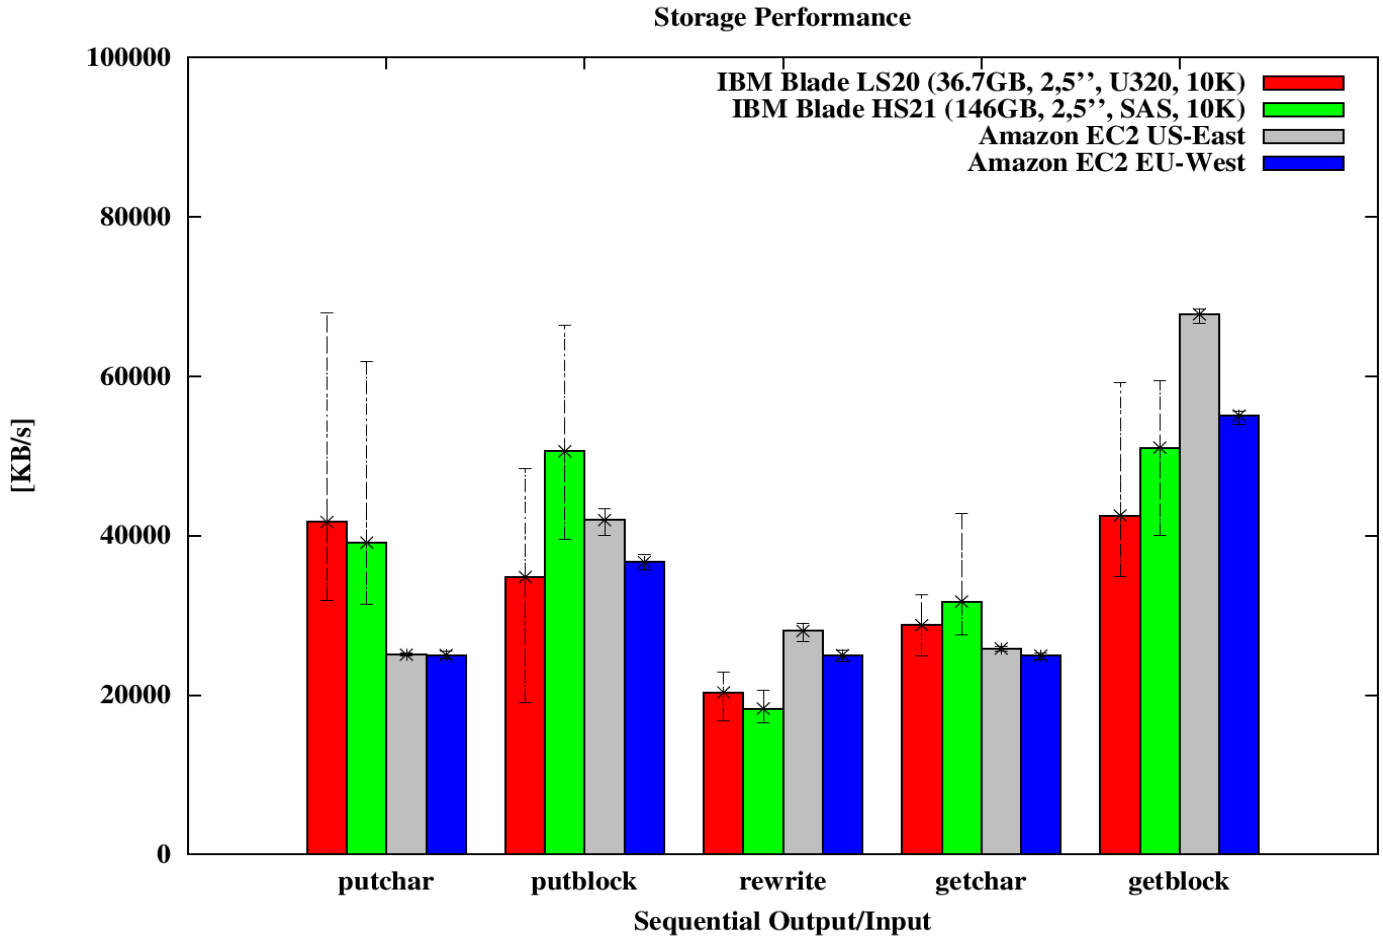

- The RAM of the Eucalyptus Node Controller was reduced to overcome the memory caching of the Linux kernel
- Reducing the RAM is not a realistic scenario but it was necessary to get correct measurements
- The write performance of Eucalyptus is better in this scenario
- Eucalyptus' storage performance depends on the available storage sub-system

Random Seeks and File Creation/Deletion

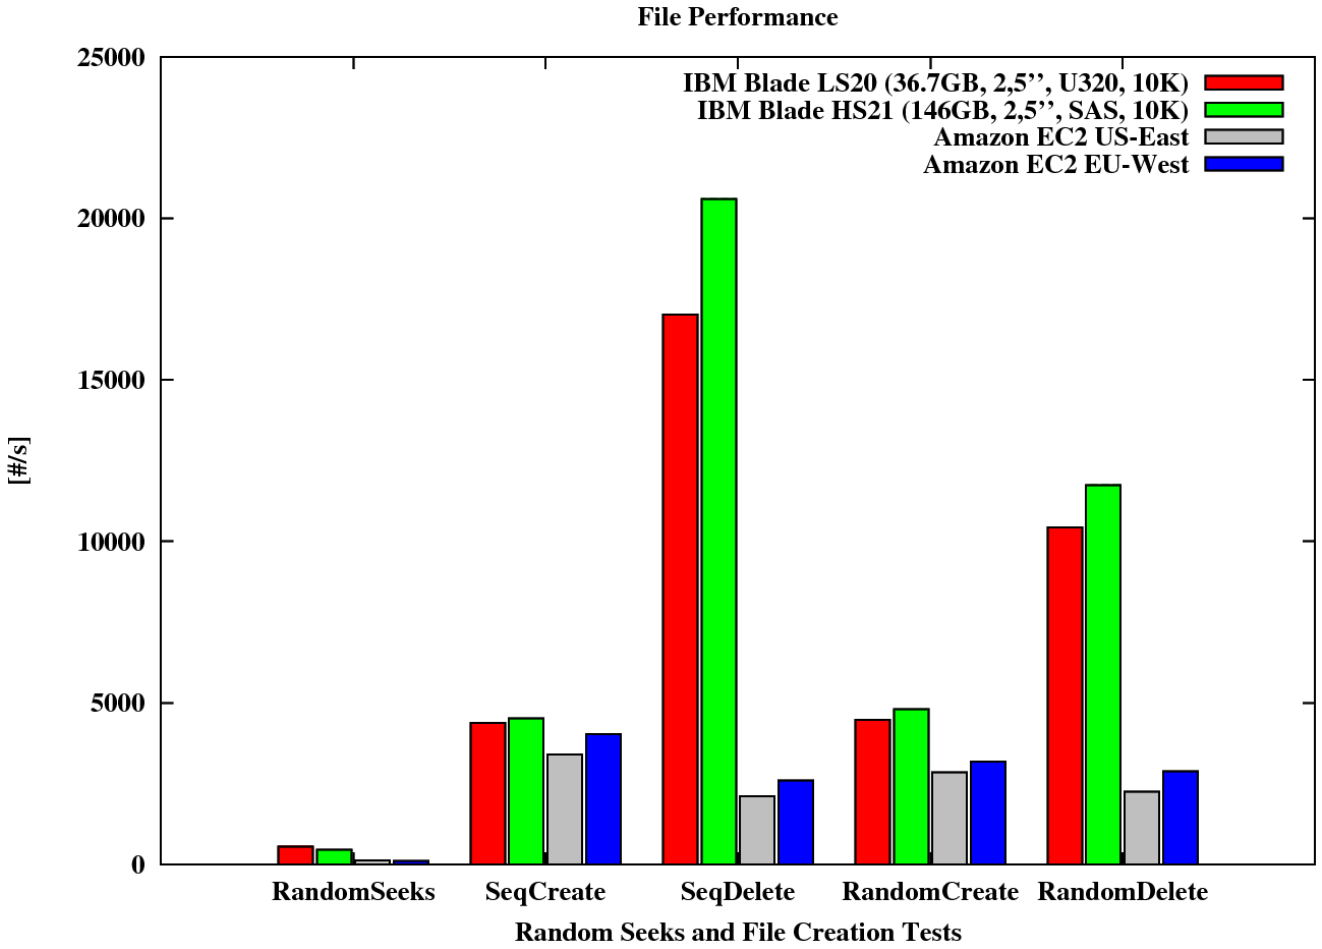

- Random seeks and file creation/deletion with Eucalyptus is faster compared to S3
- Because of the close distance between the EC2 and S3 components?
- Storage subsystem of Amazon S3 is unknown
- Performance capability and workload of Amazon S3 is unknown also

CPU Performance

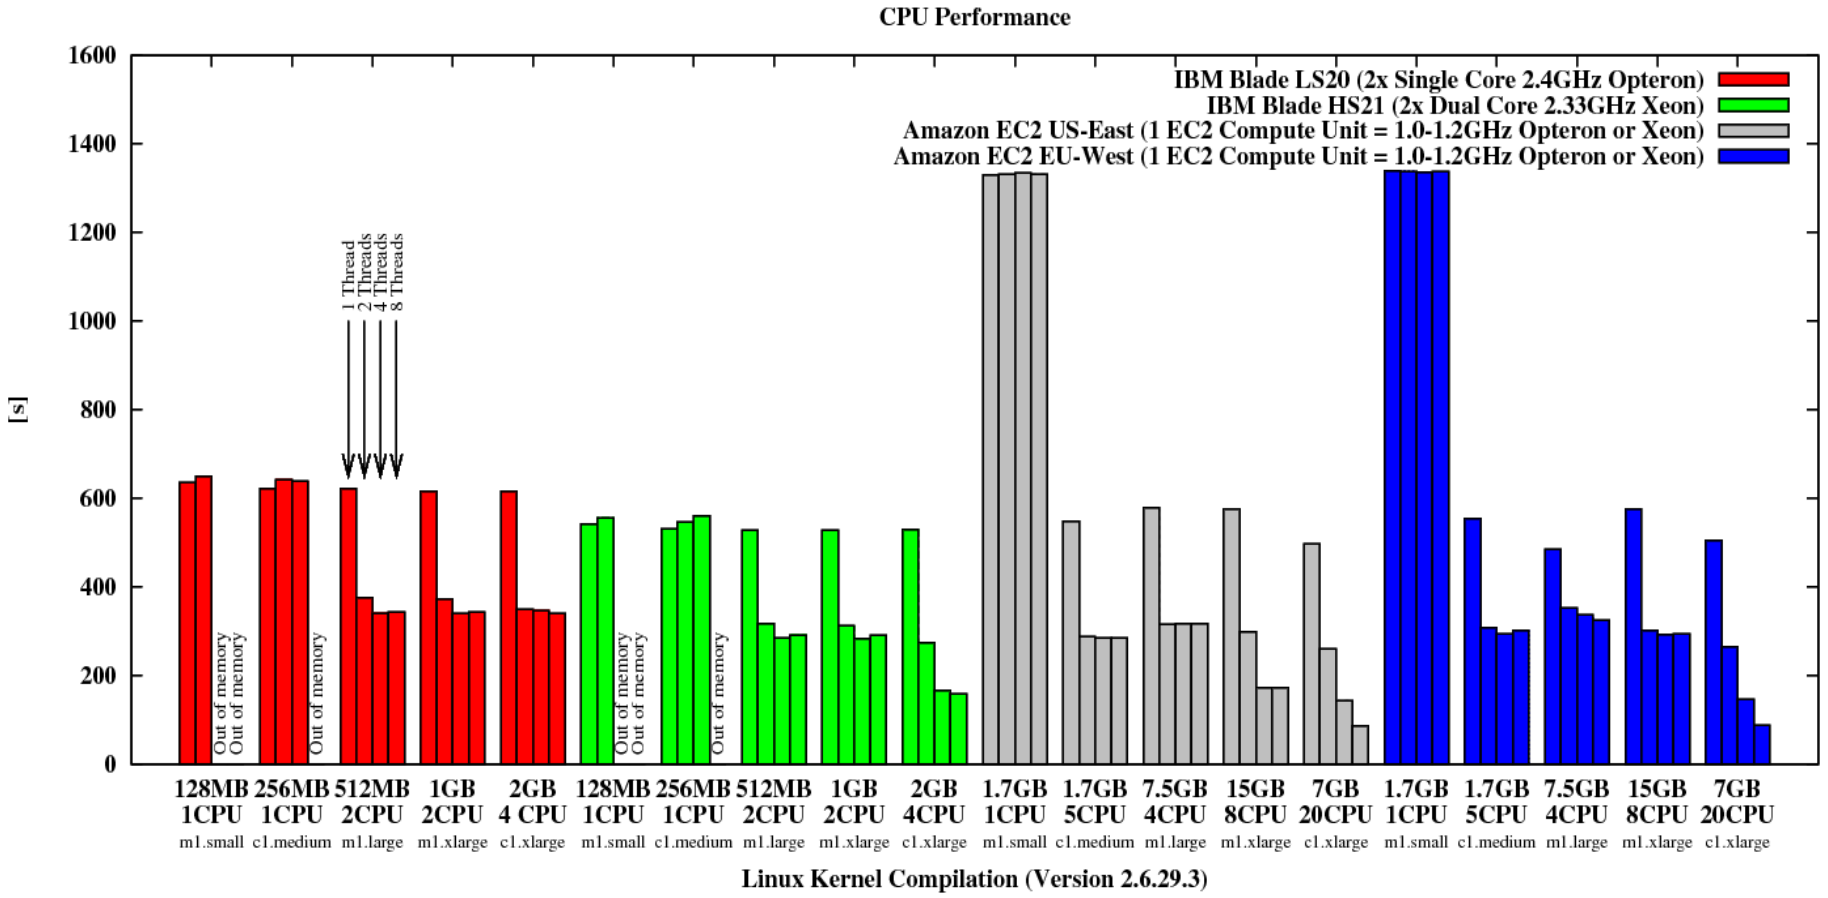

- The performance of this Eucalyptus infrastructure with the default instance types is comparable to EC2
- At Eucalyptus, every virtual CPU was mapped to one physical core. No resource sharing per default
- Using more threads than virtual/physical CPUs/cores available is not leading to a performance boost
- Using more than 2 threads at Amazon EC2 with `c1.medium` and `m1.large` is not helpful (?!)
- Using more than 4 threads at Amazon EC2 with `m1.xlarge` is not helpful (?!)

Network Transfer Rate

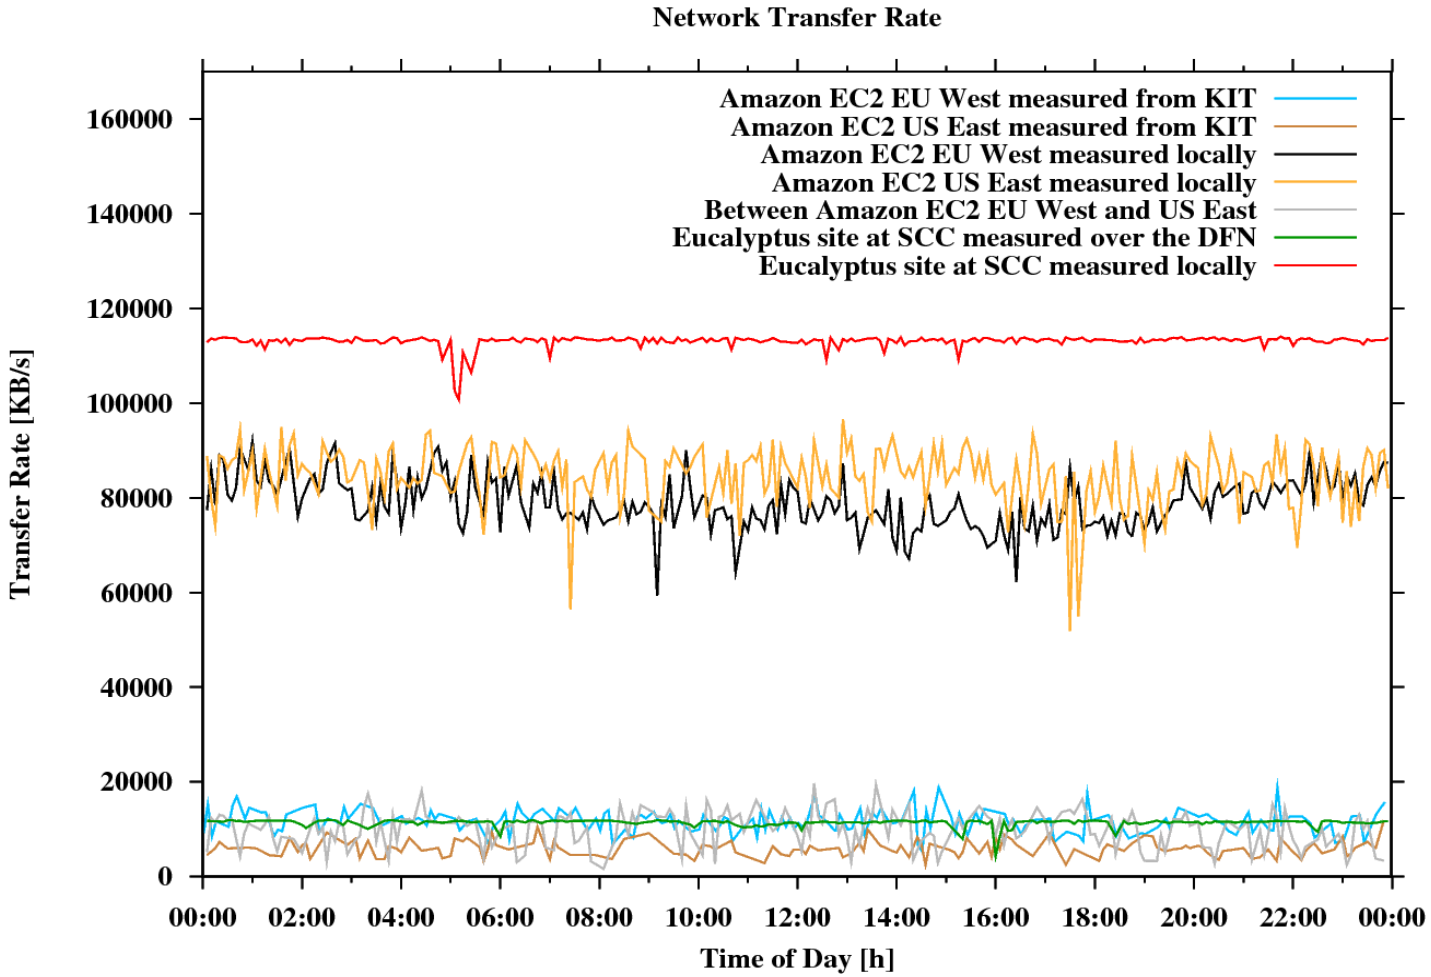

- Network Transfer Rate from/to Eucalyptus at FZK and Amazon EC2 (also measured locally)
- This graph is confusing ...

Network Transfer Rate - More in Detail (1)

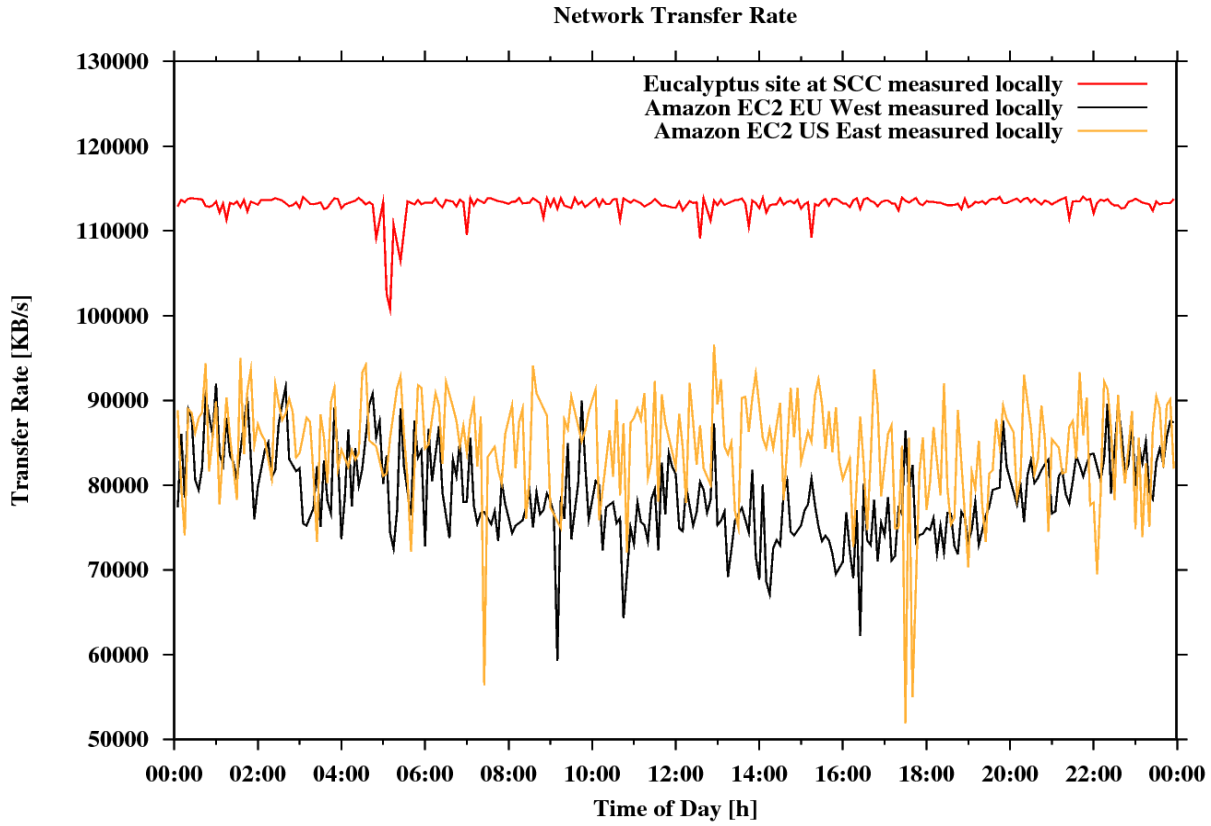

- Network Transfer Rate within the Eucalyptus site is more constant compared to Amazon EC2
- The reason for the strong Network Transfer Rate between the Eucalyptus Nodes is the 1000 Mbit/s Ethernet
- We assume Amazon uses 1000 Mbit/s Ethernet in their EC2-Sites too but there is much more workload

Network Transfer Rate - More in Detail (2)

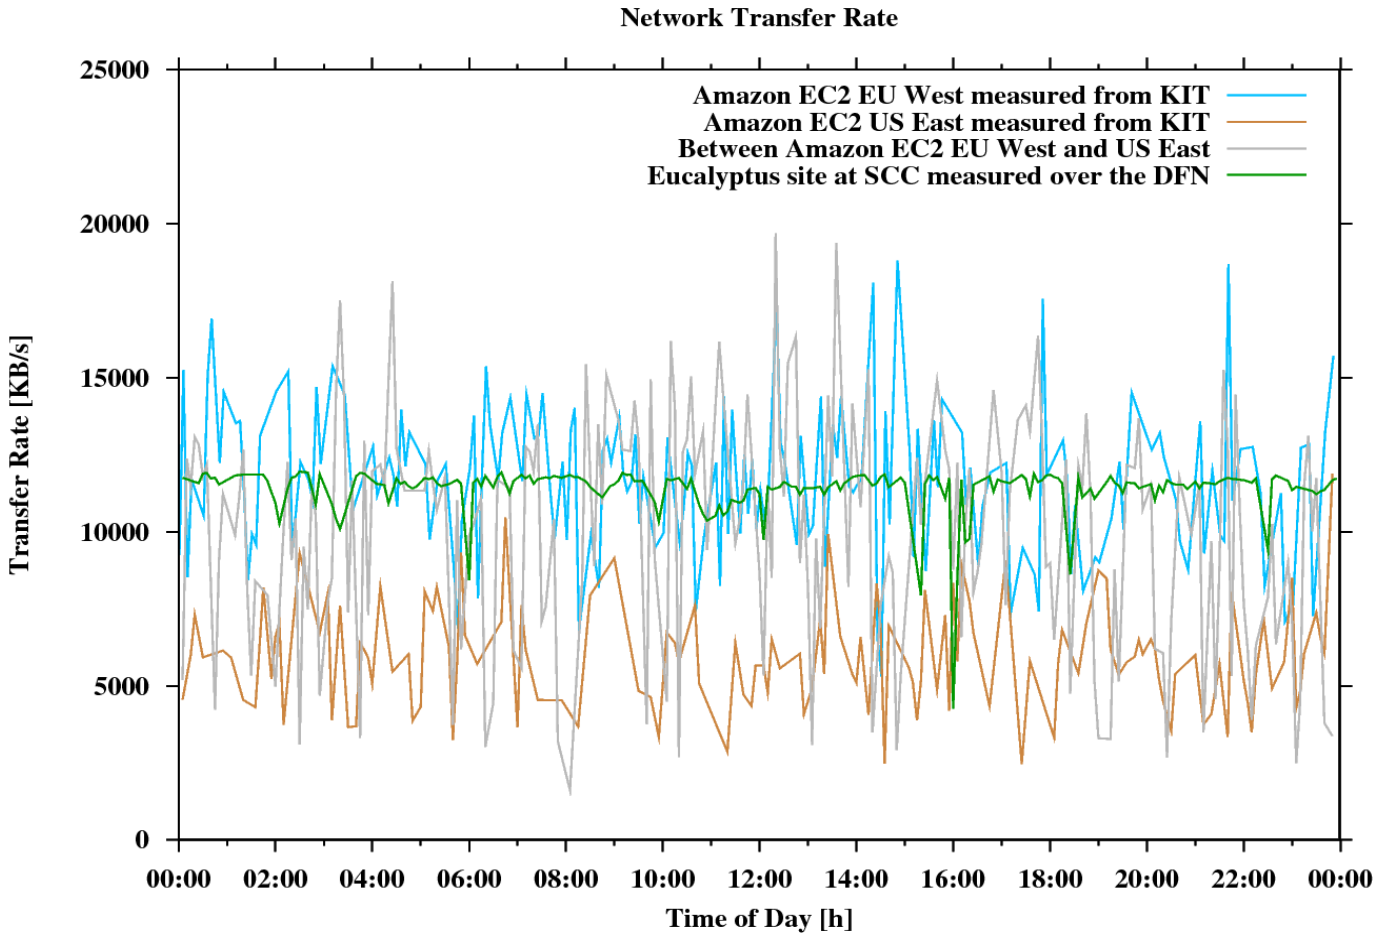

- The Network Transfer Rate from inside Europe to Amazon EC2 EU is much better compared to Amazon EC2 US. This is not surprising
- The Network Transfer Rate to Eucalyptus over DFN is more constant compared to EC2
- Strange: Peaks of the Network Transfer Rate to EC2 EU are much better compared to the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

Network Latency

ping ip -f -c 10000	time	min. Round-Trip-Time [ms]	avg. Round-Trip-Time [ms]	max. Round-Trip-Time [ms]
Amazon EC2 EU West measured from KIT	138262	27.943	28.192	59.399
Amazon EC2 US East measured from KIT	137014	92.839	93.154	118.853
Amazon EC2 EU West locally	146447	87.493	90.069	145.109
Amazon EC2 US East locally	147380	87.527	92.266	115.461
Between EC2 EU and EC2 US	138451	88.260	90.776	144.078
Eucalyptus site at SCC measured over the DFN	131145	15.093	15.197	29.863
Eucalyptus site at SCC locally	2064	0.125	0.146	0.806

- **The measurements show:**
  - **HPC in the cloud over institutional/geographical borders is impossible**
  - **For MPI-Jobs where every task computes a few seconds like Monte Carlo methods it is possible to use cloud infrastructures**
  - **Network Latency inside Amazon EC2 is very bad**
  - **Strange: Network Latency between EC2 EU West and EC2 US East is better than inside the sites**

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

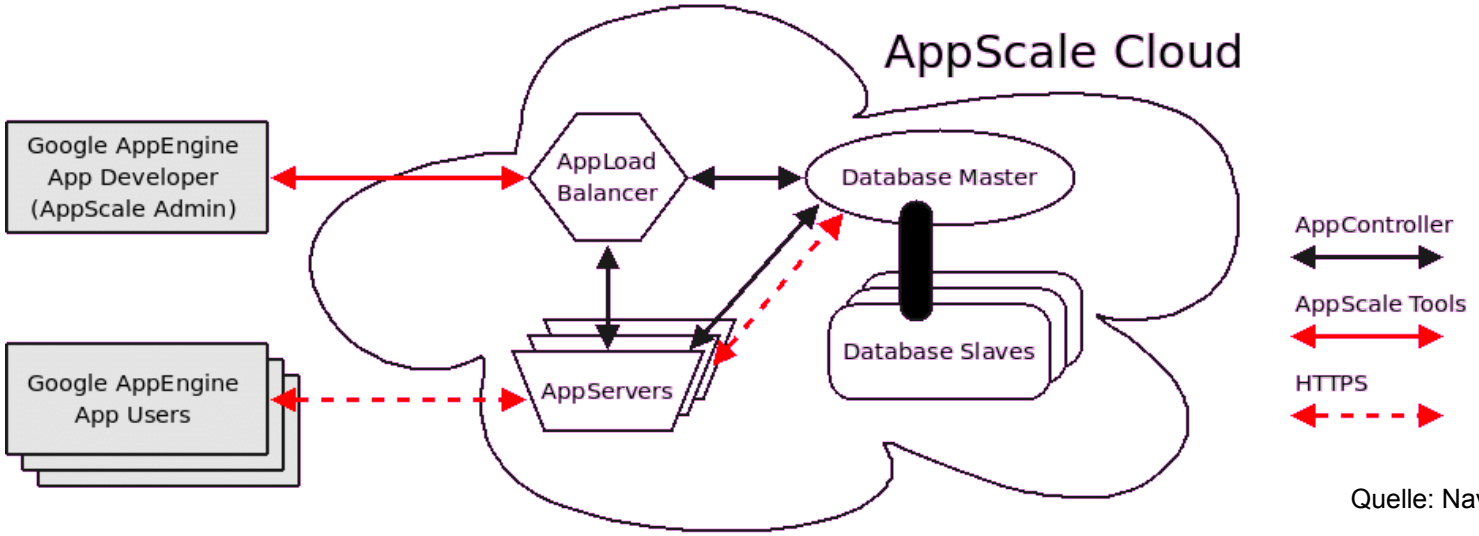

Quelle: Navraj Chohan

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 and PaaS as Open Source solutions with Linux
- With commodity hardware and Open Source software, a private cloud can be build up providing the same functionality and better performance compared to the most popular public clouds

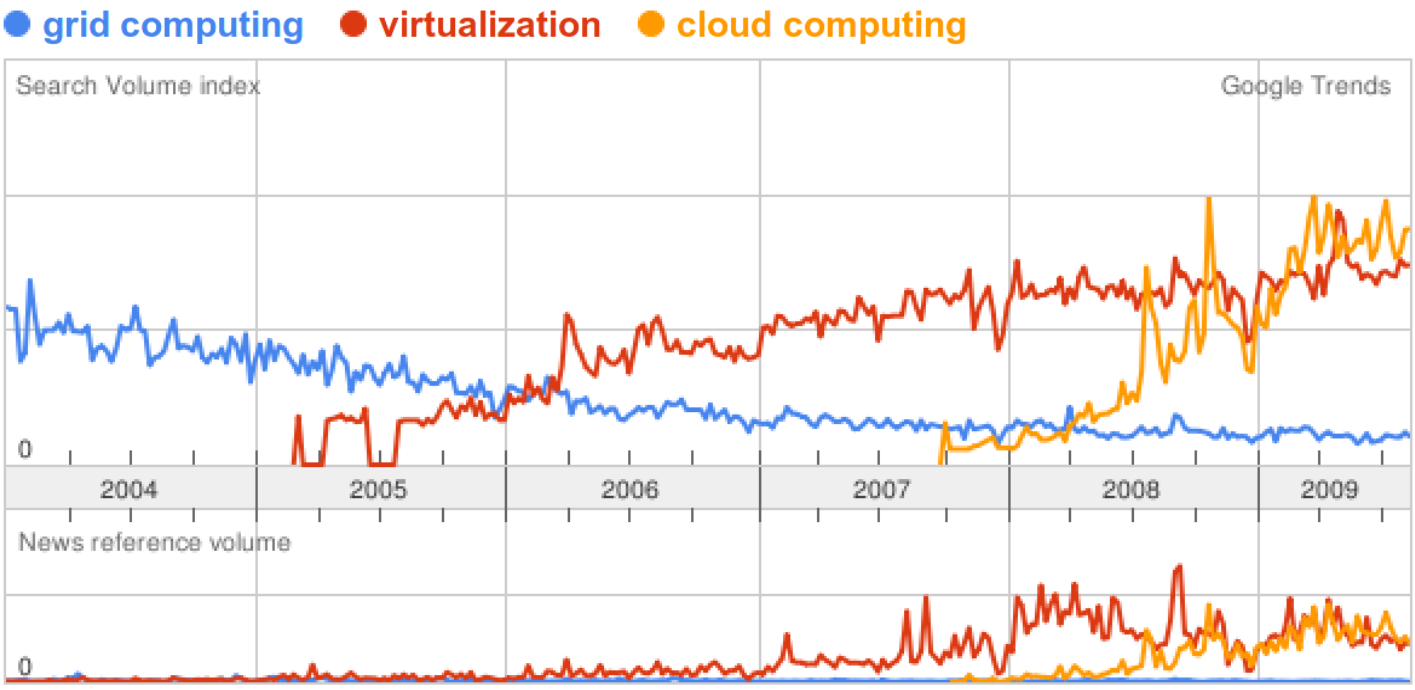

Coming soon - this autumn

- 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