Karlsruhe Institute of Technology

Assembling Cloud Infrastructures with Eucalyptus

Cloud Expo Europe 2009 (London)

Christian Baun

Forschungszentrum Karlsruhe
in der Helmholtz-Gemeinschaft

Universität Karlsruhe (TH)
Forschungsuniversität • gegründet 1825

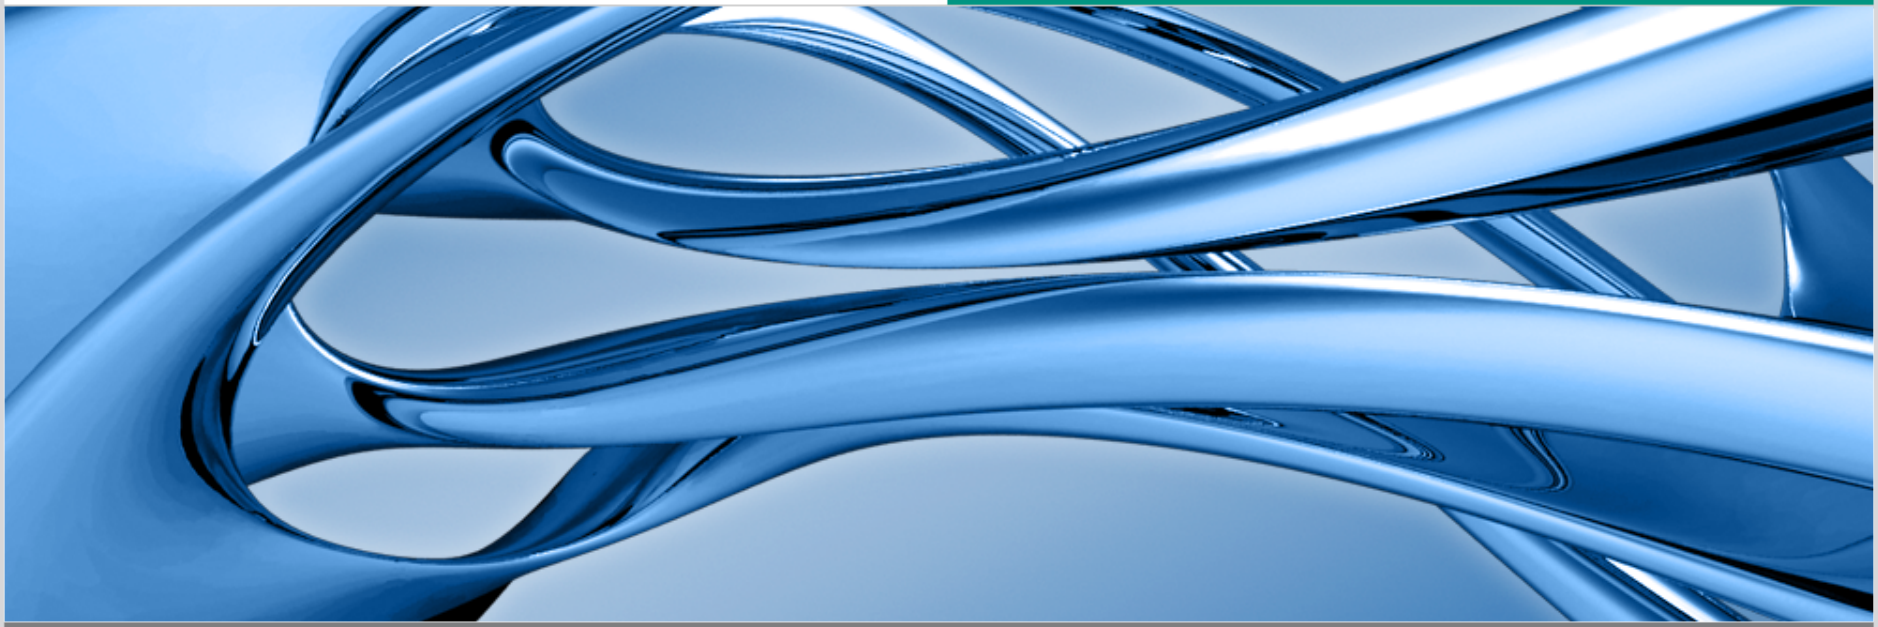

Agenda

- Clouds vs. Grids
- Popular Cloud/Grid Infrastructures
- Types of Cloud Services
- The OpenCirrus™ project
- Eucalyptus
- AppScale

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

Clouds vs. Grids: A Comparison

	Cloud Computing	Grid Computing
Objective	Provide desired computing platform via network enabled services	Resource sharing Job execution
Infrastructure	One or few data centers, heterogeneous/homogeneous resource under central control Industry and Business	Geographically distributed, heterogeneous resource, no central control, VO Research and academic organization
Application	Suited for generic applications	Special application domains like High Energy Physics
Business Model	Commercial: Pay-as-you-go	Publicly funded: Use for free (negotiate with resource owner for resource access)
Middleware	Proprietary, several reference implementations exist (e.g. Amazon)	Well developed, maintained and documented
User interface	Easy to use/deploy, no complex user interface required	Difficult use and deployment Need new user interface, e.g., commands, APIs, SDKs, services …
Operational Model	Industrialization of IT Fully automated Services	Mostly Manufacture Handcrafted Services
QoS	Possible	Little support
On-demand provisioning	Yes	No

THE Grid-Project: LHC Computing Grid (1)

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

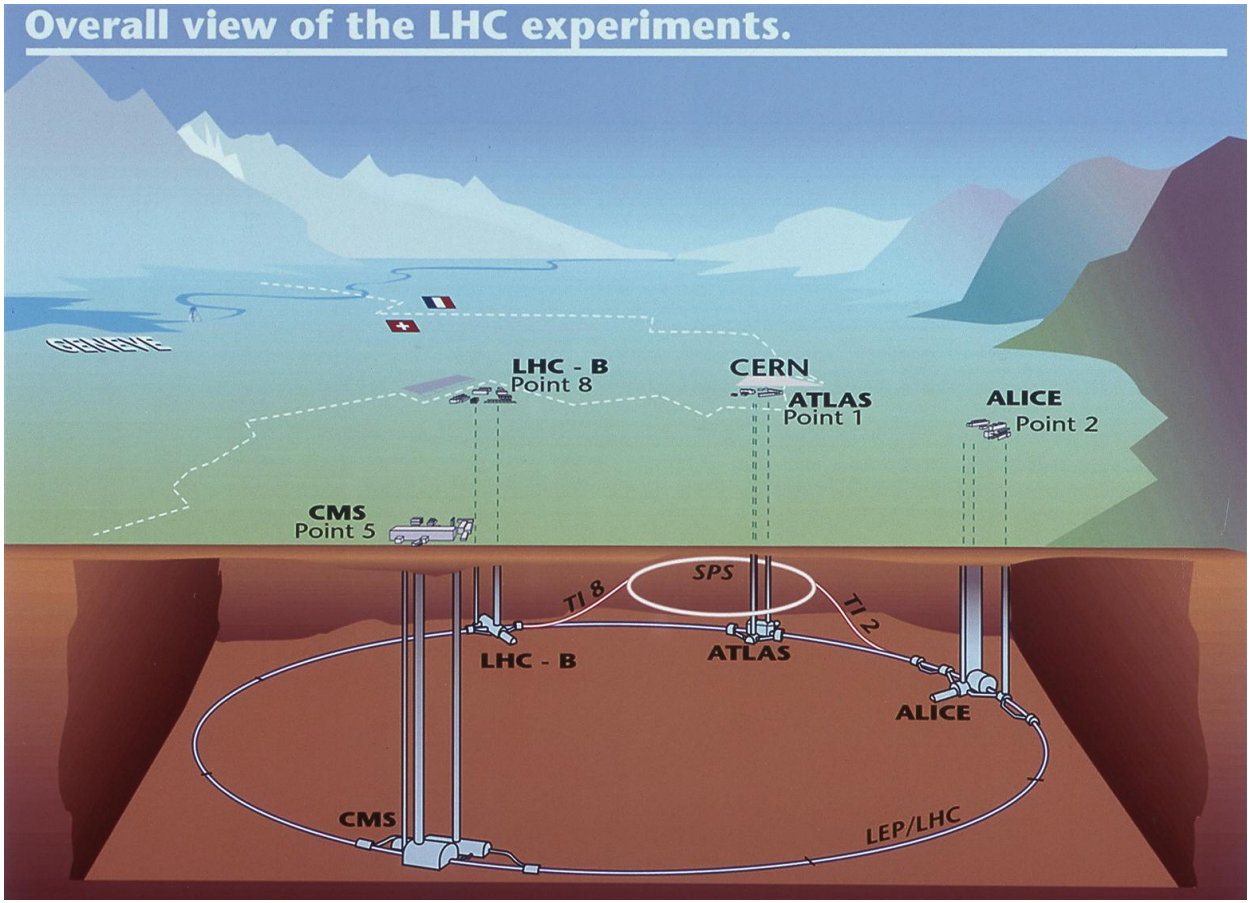

Source: http://guillegg.wordpress.com

- 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/

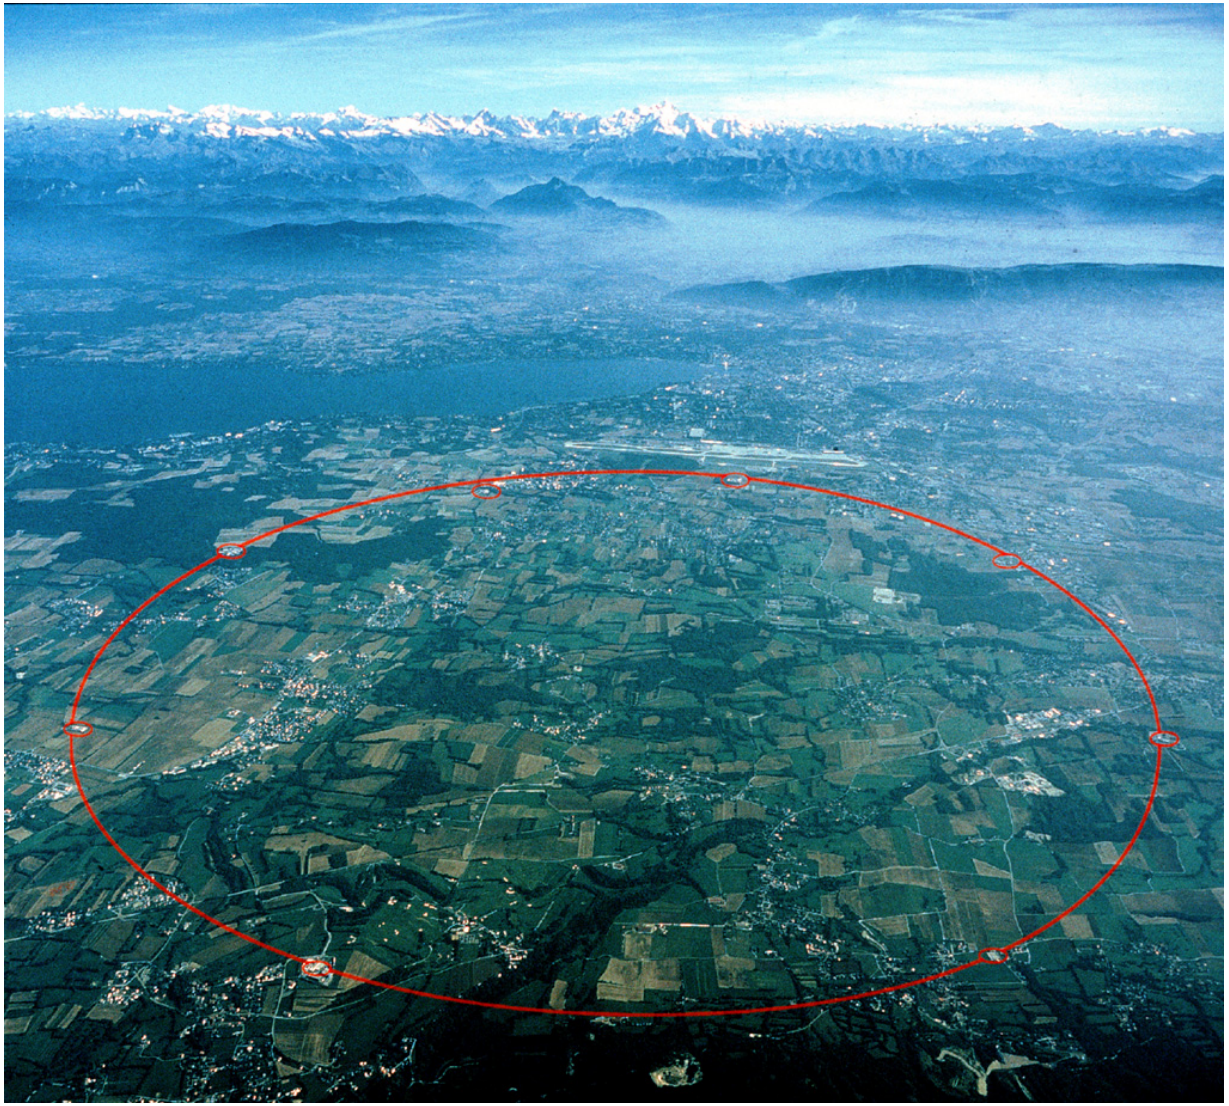

- The LHC generates 10 - 40 Petabyte of data for all experiments per year
- Data is stored and analyzed distributed inside the LHC-Grid
  - Largest computing grid in the world
  - > 150 Sites
  - > 20000 Servers

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

EGEE: 60000+ CPUs with Linux

http://gridportal.hep.ph.ic.ac.uk/rtm/

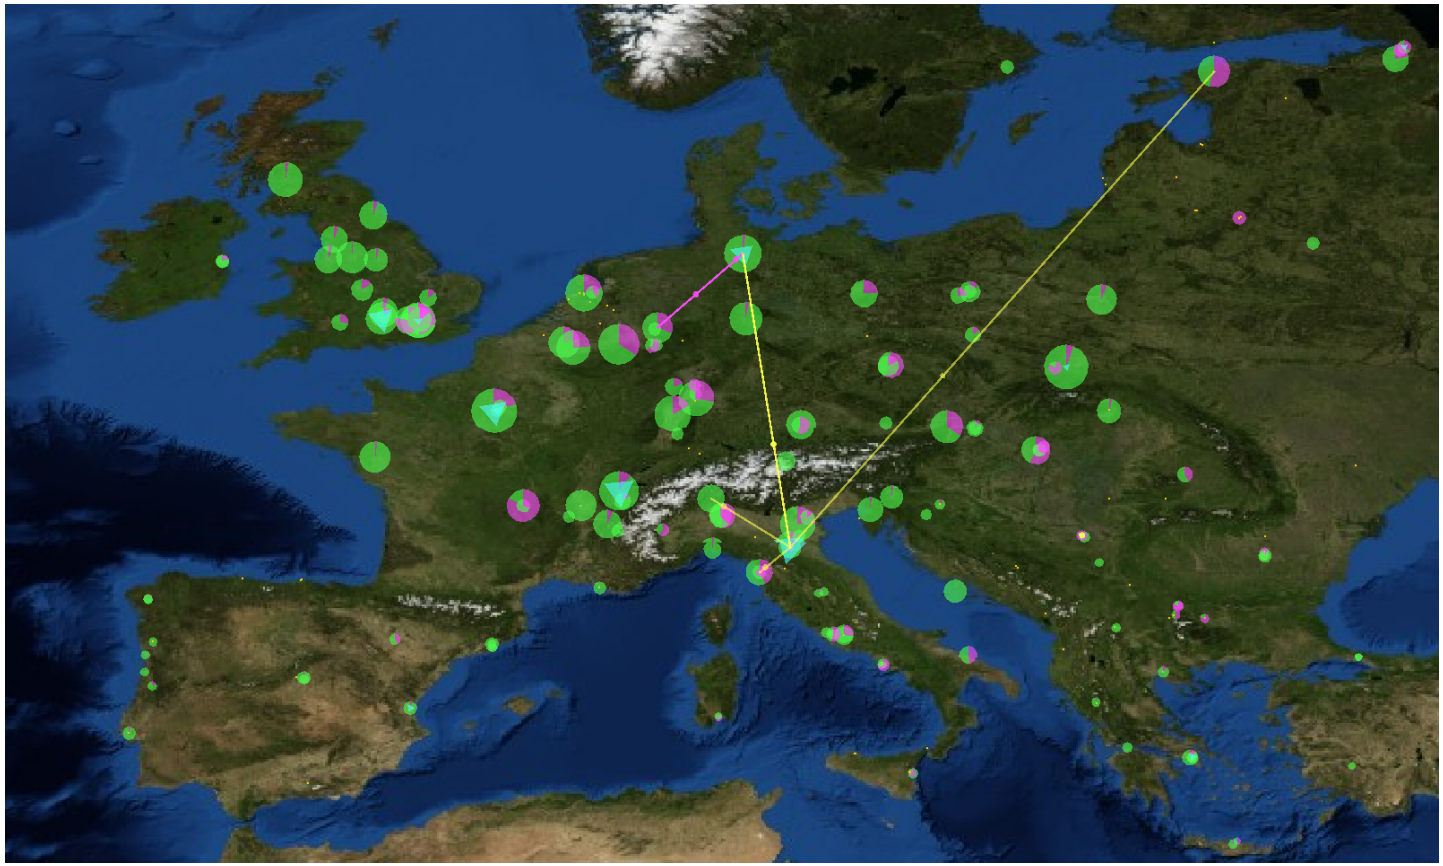

- **The Real Time Monitor shows:**
  - **running and scheduled jobs**
  - **job transfers**
  - **detailed information on Resource Brokers and Computing Elements for each site**

Three Major Types of Cloud Services

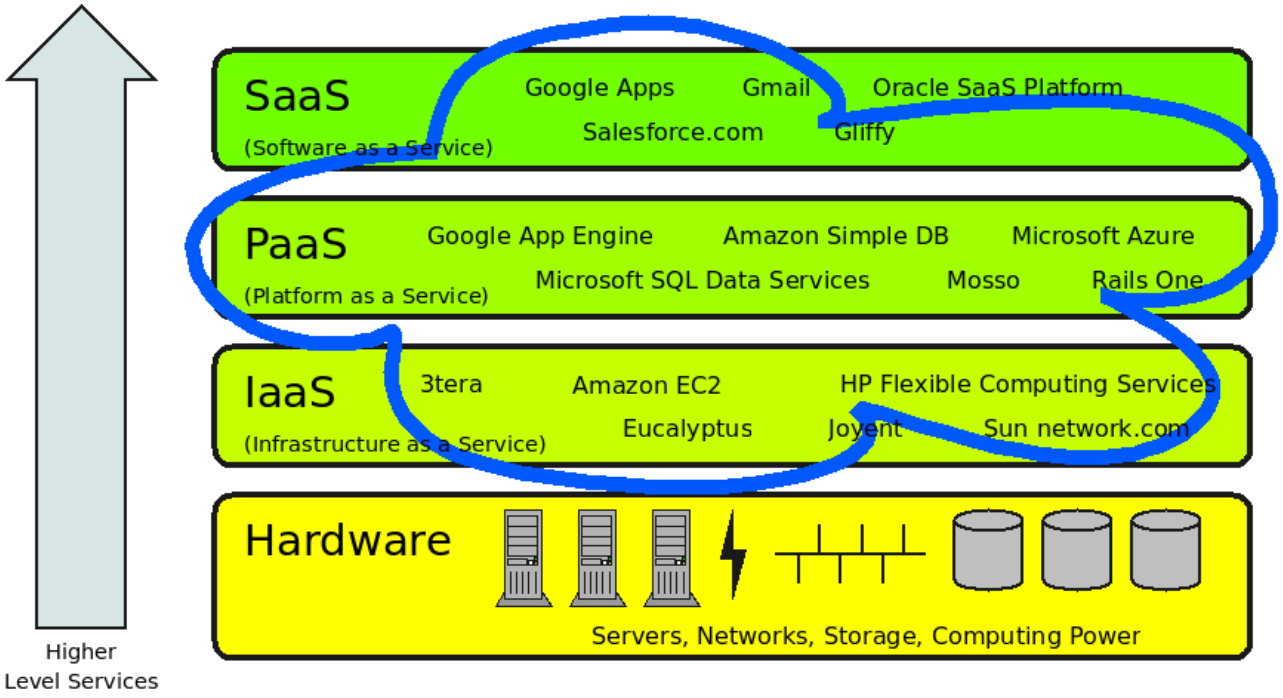

- **SaaS:**
  - **Provides enterprise quality software (complete applications)**
- **PaaS:**
  - **Appears as one single large computer and makes it simple to scale from a single server to many**
  - **No need to worry about the operating system or other foundational software**
- **IaaS:**
  - **Abstracts away the hardware (servers, network,…) and allows to run virtual instances of servers without ever touching a piece of the hardware**

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 Summer 2009
  - 3300 Nehalem cores, 10TB memory, 192TB hard disk storage

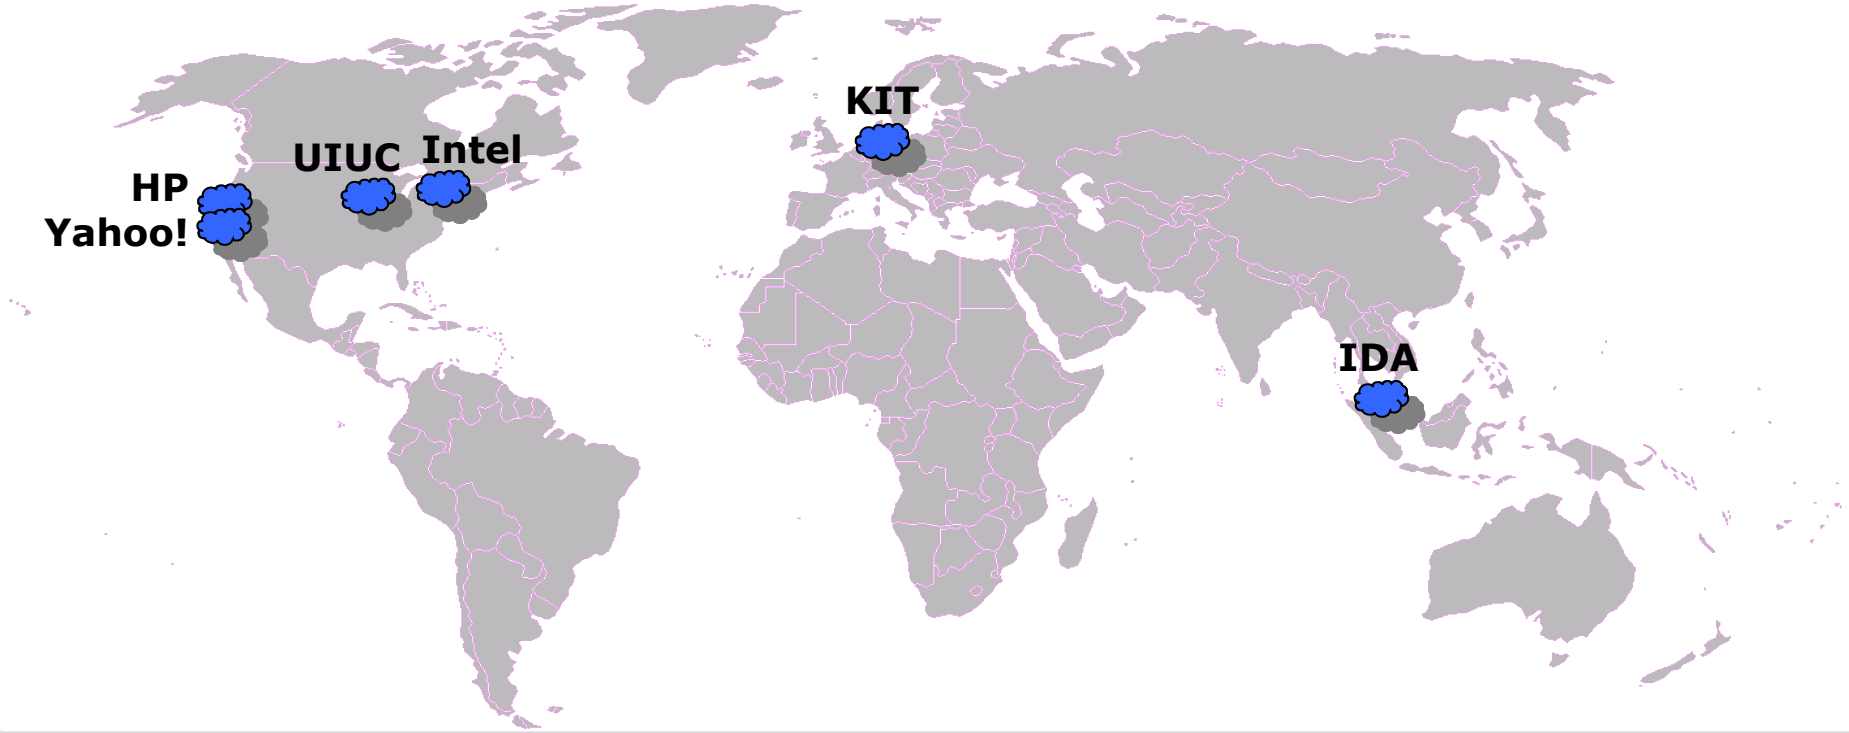

OpenCirrus™ - Physical Resource Sets (PRS)

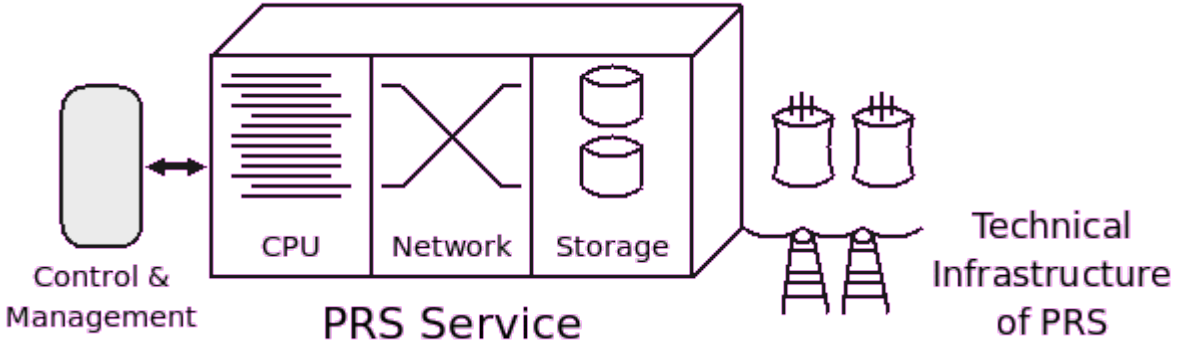

- PRS service goals
  - Provide mini-datacenters to the researchers
  - Isolate experiments from each other
- PRS service approach
  - Allocate sets of physical co-located nodes, isolated inside VLANs using existing software
    - Utah Emulab - Network Emulation Testbed
    - HP Opsware - Server provisioning, configuration and management
    - …
  - Start simple, add features as we go
  - Basis to implement Virtual Resource Sets (VRS)
- Hardware as a Service (HaaS)

OpenCirrus™ - Virtual Resource Sets (VRS)

- Basic idea: Abstract from physical resources by the introduction of a virtualization layer
- Concept applies to all IT aspects: CPU, storage, networks and applications, …
- Main advantages
  - Implement IT services exactly fitting customers varying needs
  - Deploy IT services on demand
  - Automated resource management
  - Easily guarantee service levels
  - Live migration of services
  - Reduce both: Capital Expenditures and Operational Expenditures
- Infrastructure as a Service (IaaS)
  - Implement Compute and Storage Services
  - De-facto standard: Amazon Web Services interface

OpenCirrus™ Blueprint

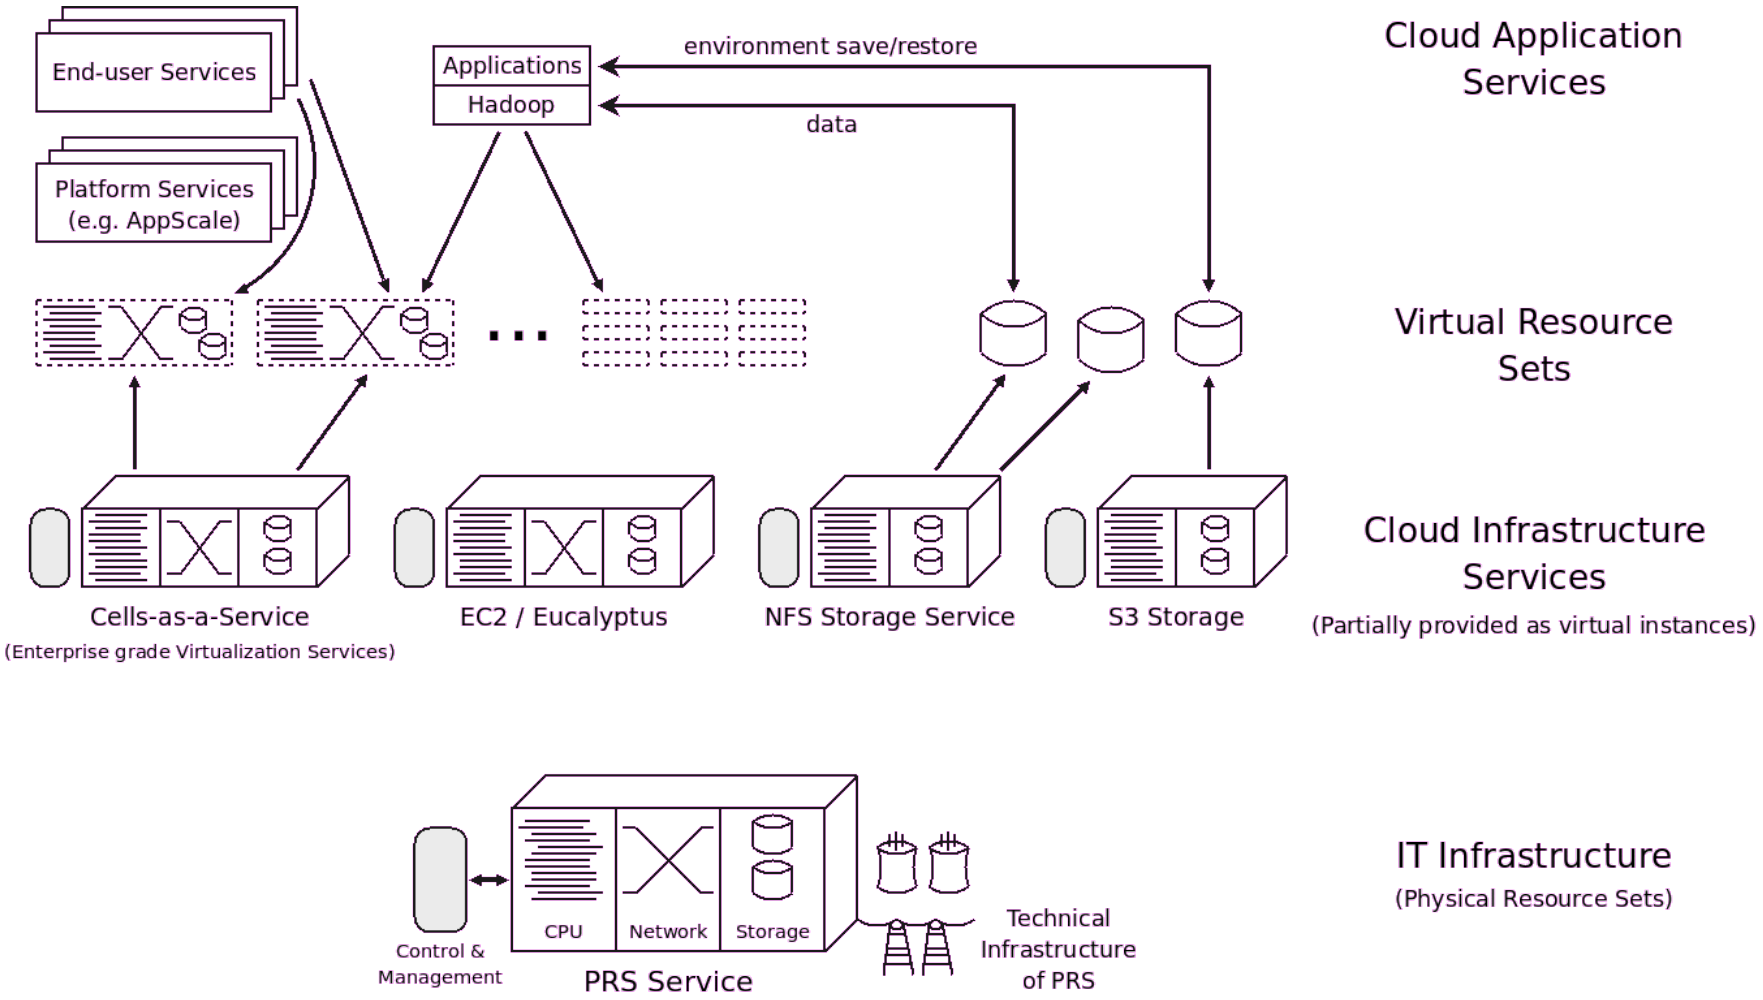

Commercial Cloud Offerings (Small Excerpt)

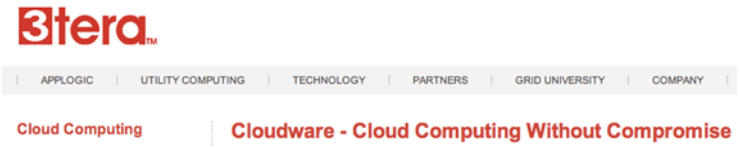

- **Problem: Commercial offers are proprietary and usually not open for Cloud systems research and development!**

Eucalyptus

http://eucalyptus.cs.ucsb.edu

- Open-Source software infrastructure for implementing Cloud computing on clusters from UC Santa Barbara
- EUCALYPTUS - Elastic Utility Computing Architecture for Linking Your Programs To Useful Systems
- Implements Infrastructure as a Service (IaaS) – gives the user the ability to run and control entire virtual machine instances (Xen, KVM) deployed across a variety of physical resources
- Interface compatible with Amazon EC2
- Includes Walrus, a storage service that is interface compatible with Amazons S3
- 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

http://eucalyptus.cs.ucsb.edu

Eucalyptus .cs.ucsb.edu

Collects resource information from the CC. Operates like a meta-scheduler in the Cloud.

Schedules the distribution of virtual machines to the NC. Collects (free) resource information.

Runs on every node in the Cloud. Xen-Hypervisor running. Provides Information about free resources to the CC.

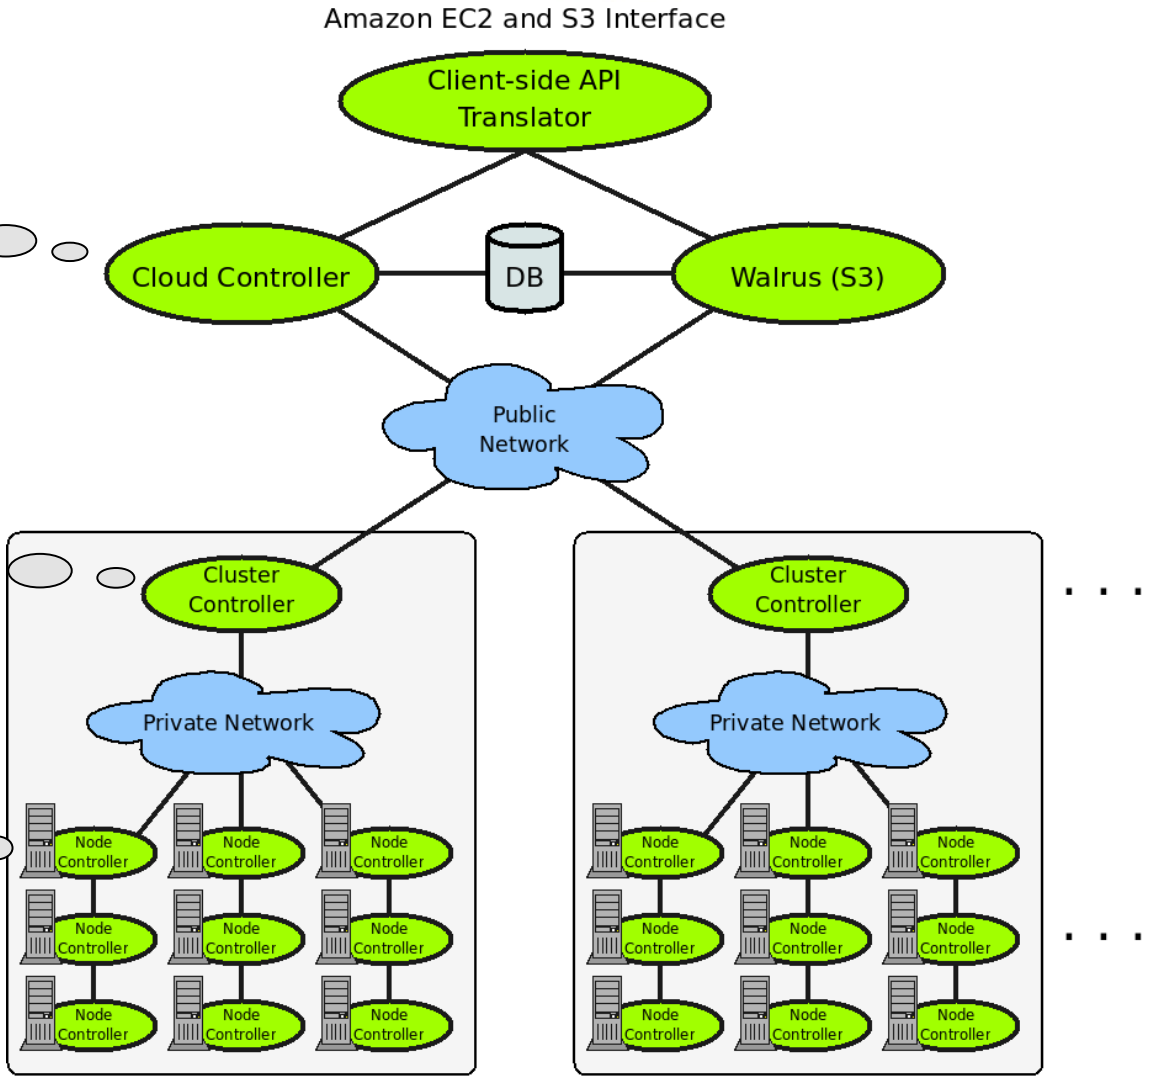

Source: R.Wolski

Comparing Storage Performance between S3 and Eucalyptus

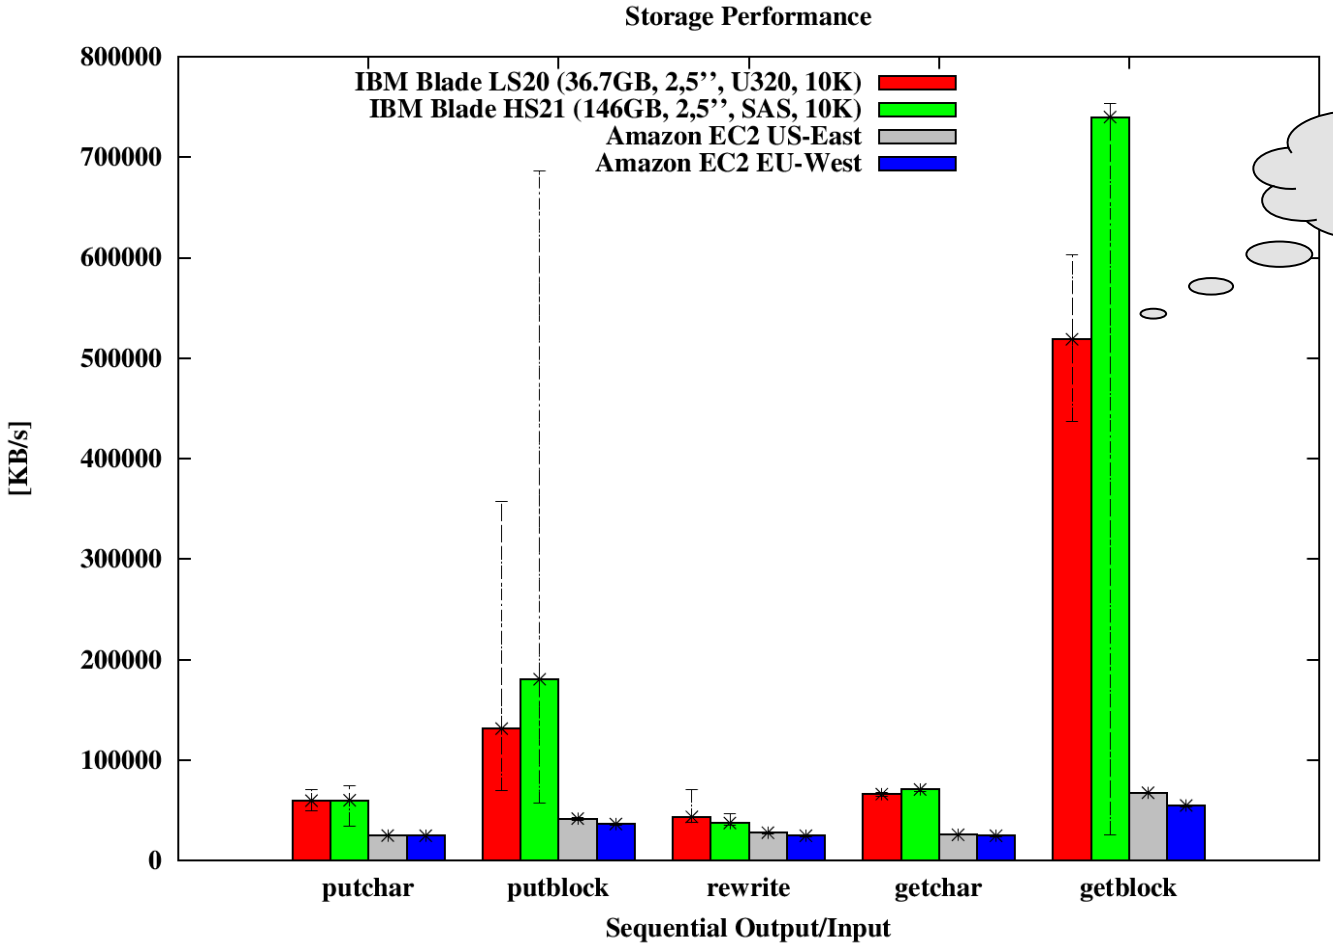

WOW!?

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

Realistic values…

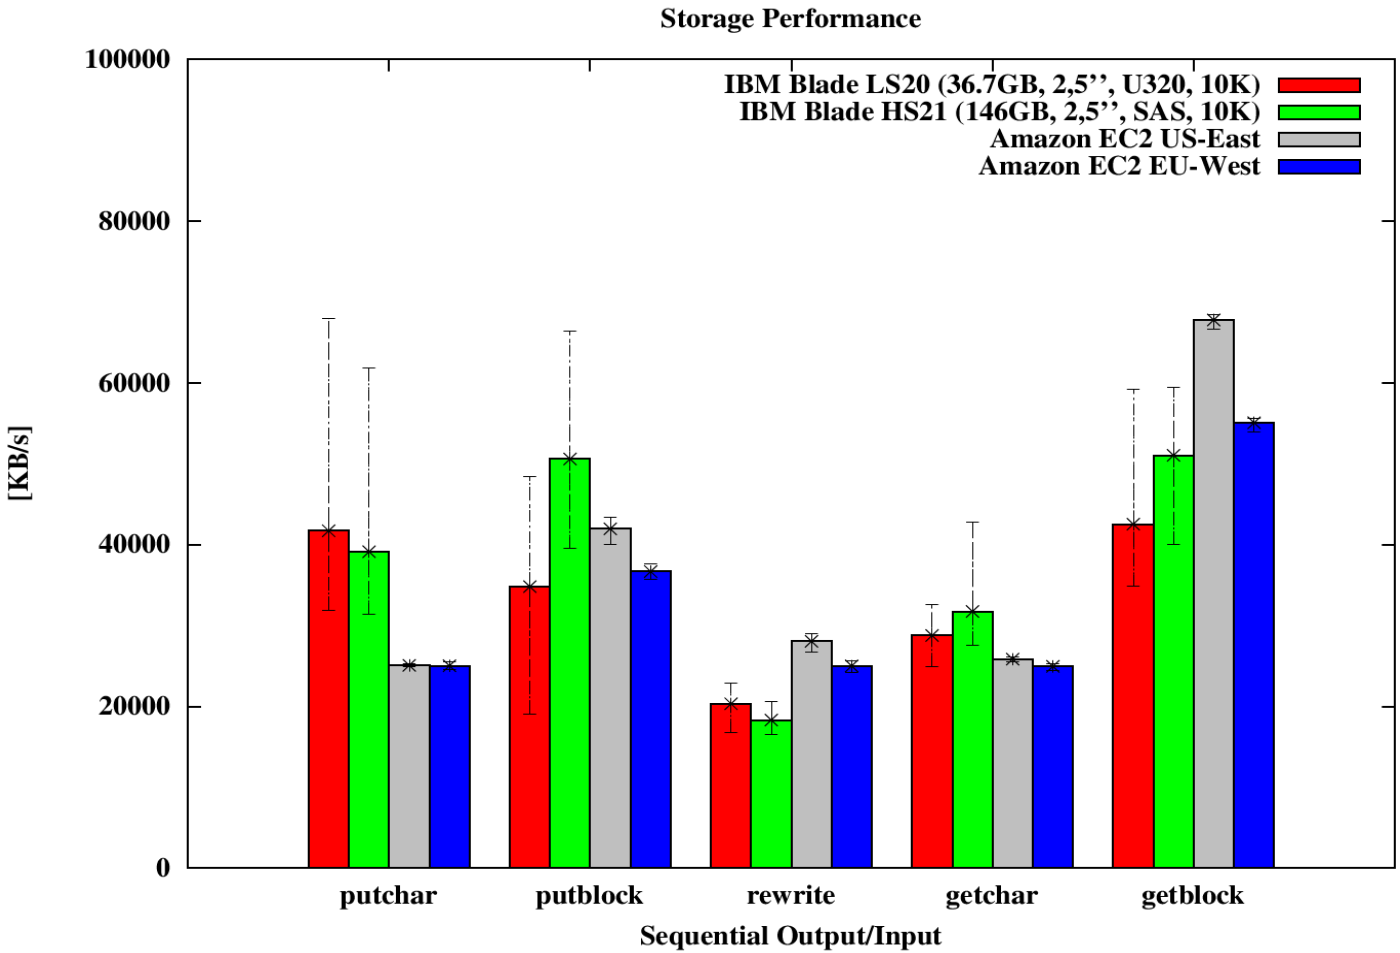

- The RAM of the Eucalyptus Node Controller was reduced to overcome memory caching
- The storage performance of Eucalyptus depends on the available storage sub-system
  - Write performance of Eucalyptus is faster. Because of the close distance?!

Performance of Random Seeks and File Creation

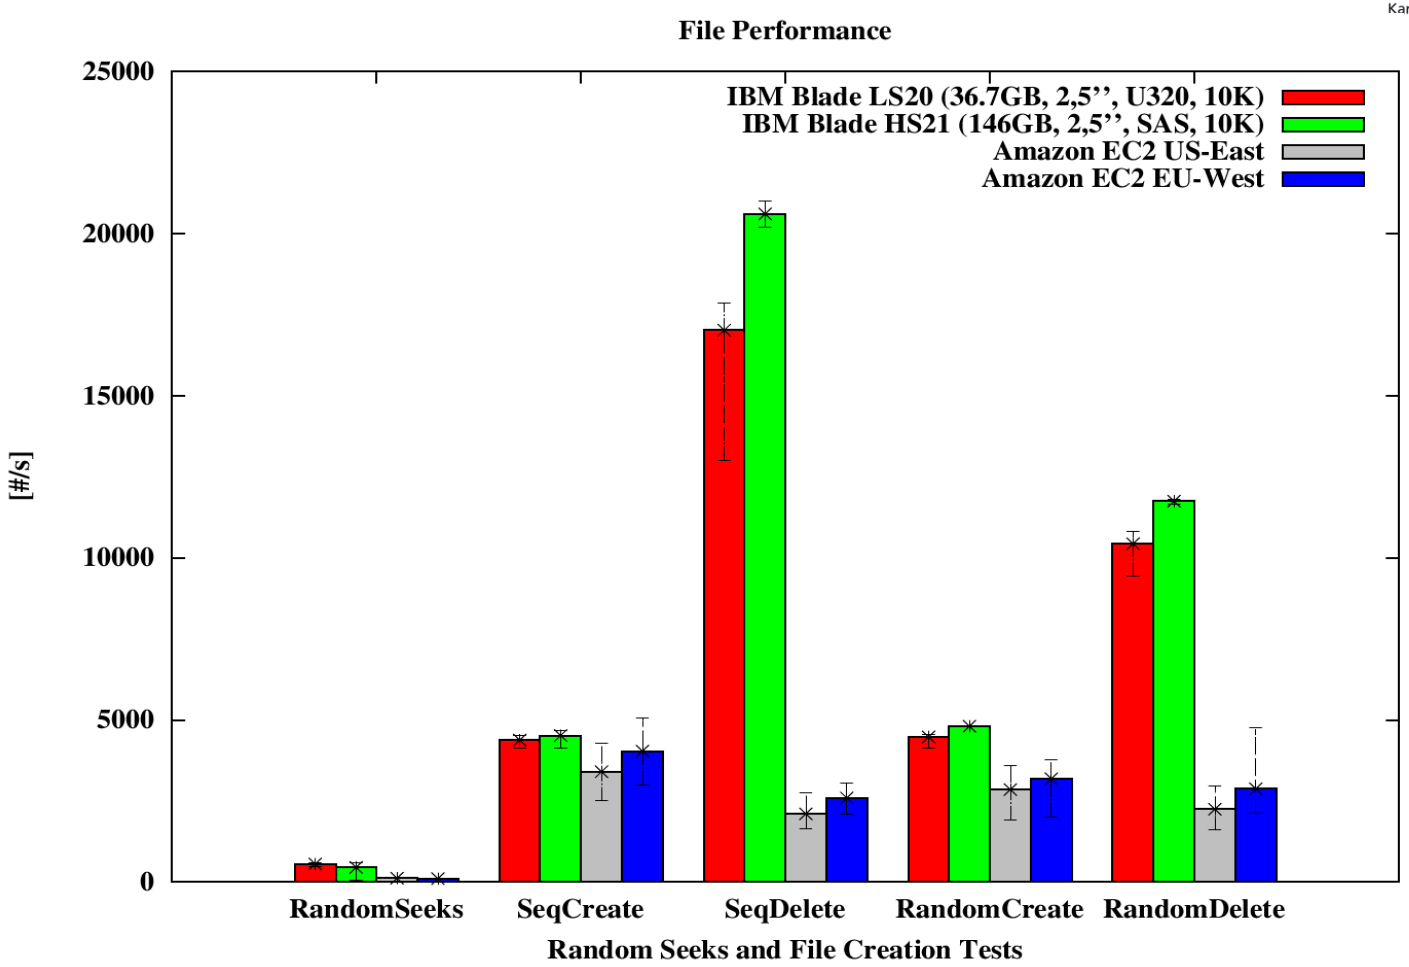

- Random seeks and file creation with Eucalyptus is faster
  - Because of the close distance?!

CPU Performance

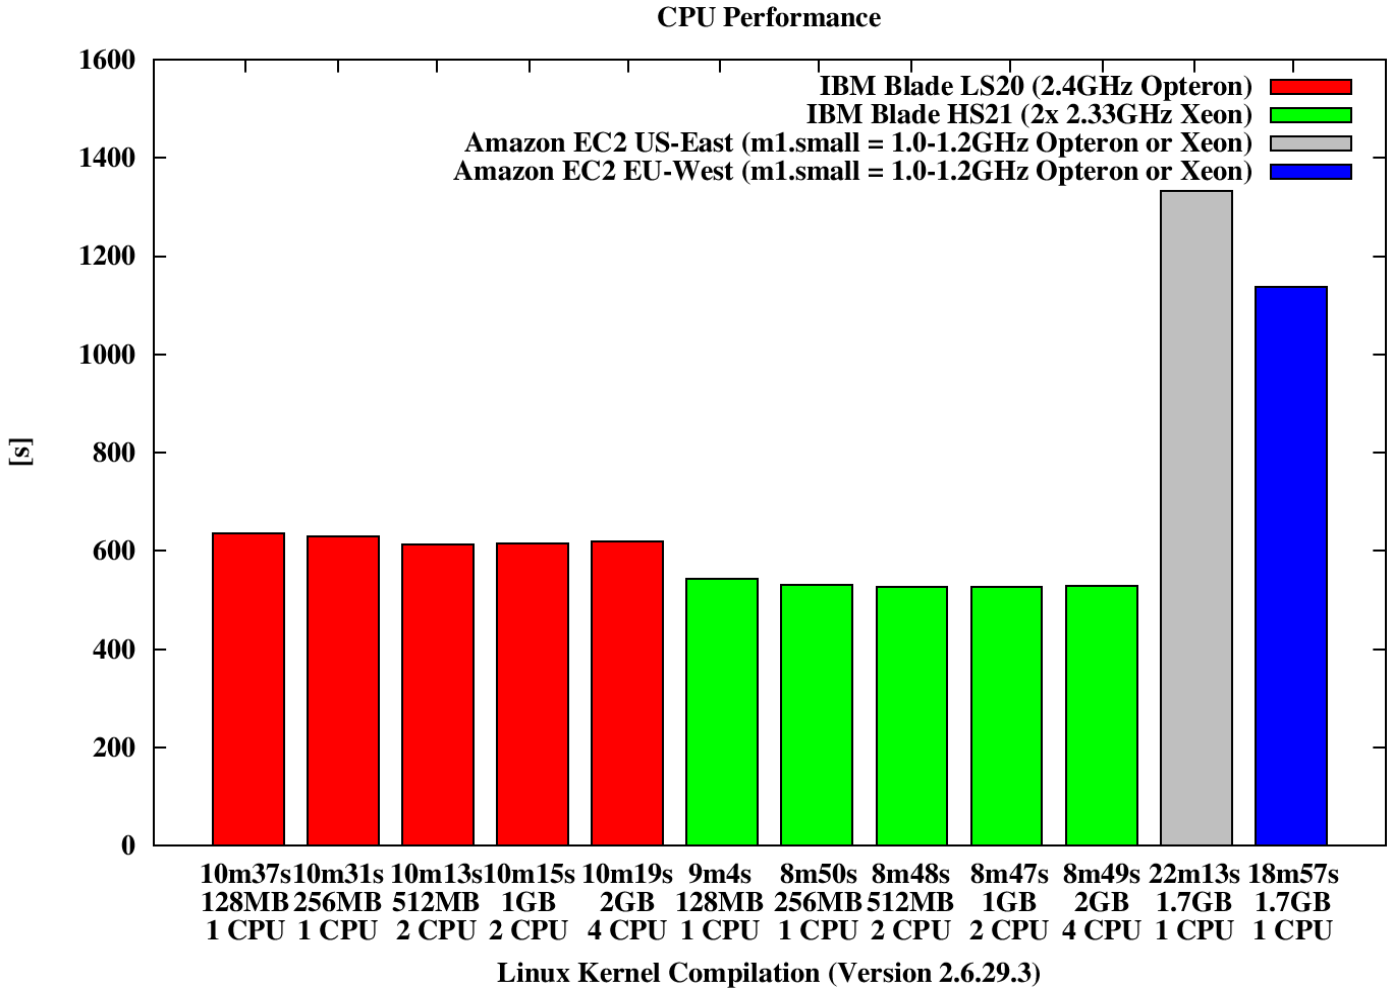

- The Eucalyptus infrastructure performs much better than Amazon EC2
  - More RAM and CPU power is not leading to a massive performance boost

Network Transfer Rate

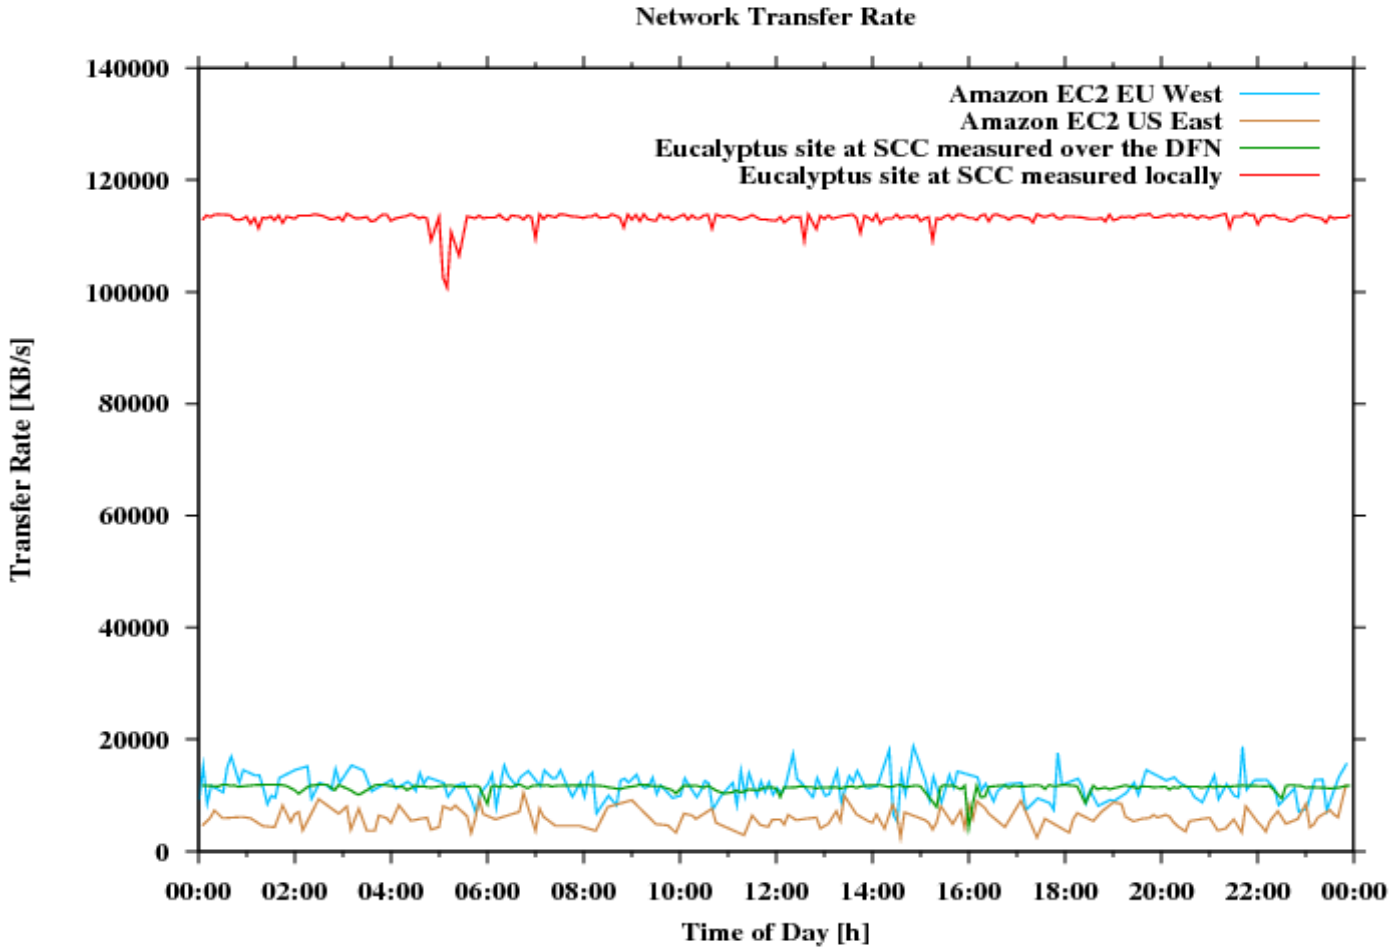

- The network transfer rate to the Eucalyptus infrastructure is more constant in contrast to Amazon EC2

Network Transfer Rate (More in Detail)

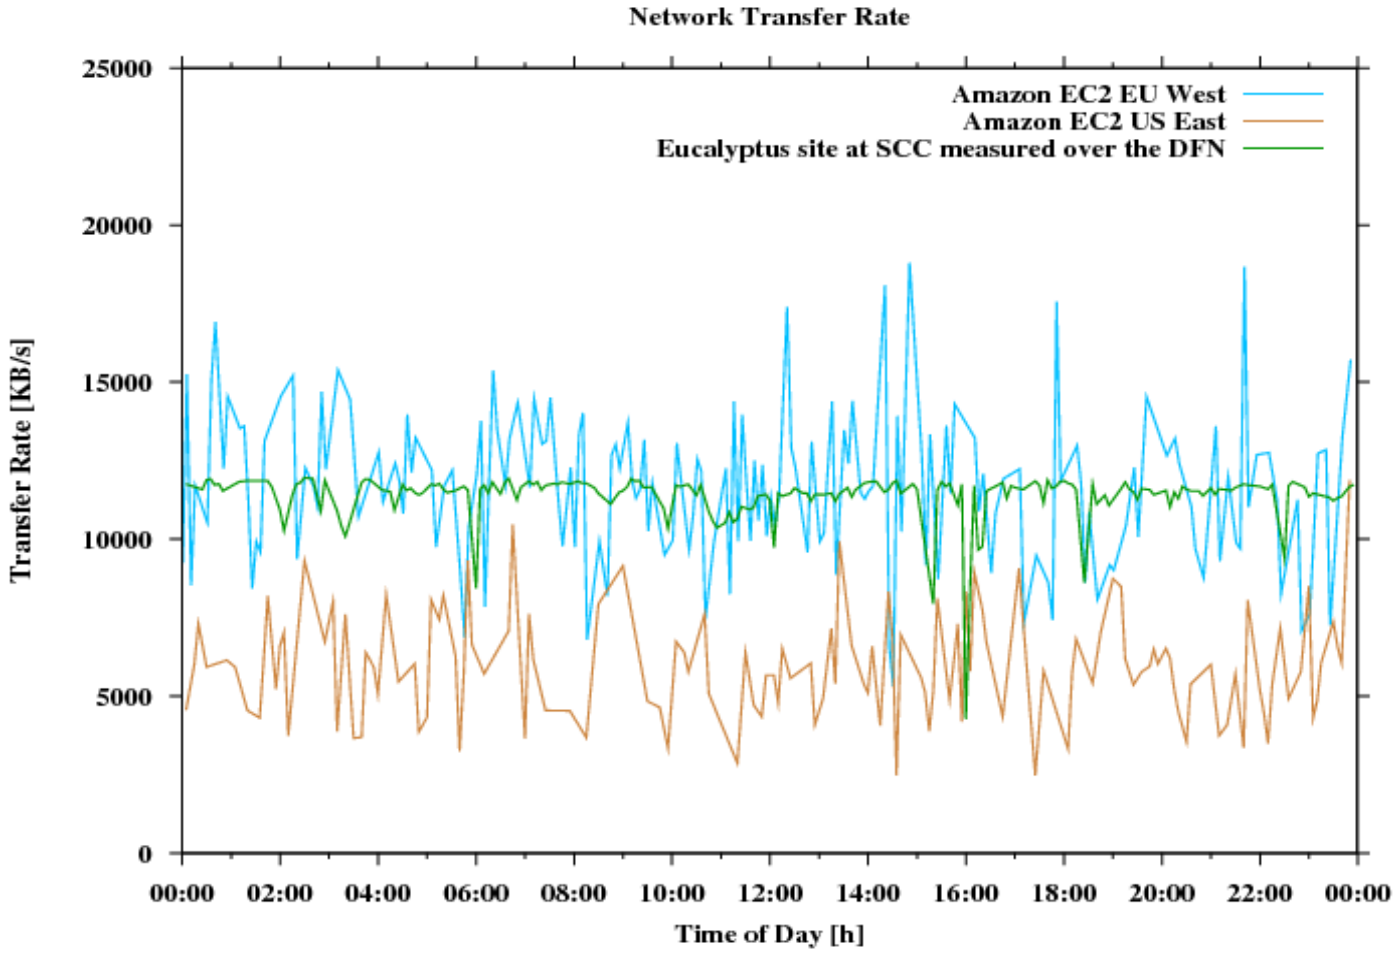

- The network transfer rate to the Eucalyptus infrastructure is more constant in contrast to Amazon EC2

Install Eucalyptus

- Installation is quite easy when using binary packackes for CentOS, openSUSE, Debian and Ubuntu

  http://open.eucalyptus.com/wiki/EucalyptusAdministratorGuide_v1.5
- Possible to build from sources

  http://open.eucalyptus.com/wiki/EucalyptusSourceCodeInstallation_v1.5
- Requirement: ≥ 1 Linux machines with working Xen Hypervisor or Kernel-based Virtual Machine (KVM)
  - For using KVM a *modern* CPU with AMD-V (Pacifica) or Intel VT (Vanderpool) is needed
- Amazon EC2 Command-Line Tools
  - ec2-api-tools-1.3-30349
  - ec2-ami-tools-1.3-26357

Use Eucalyptus

- **Shows the cluster's front-end hostname, free resources, 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
```

Run Instances

- **Run instances:**

```
# ec2-run-instances emi-1DE4116D --kernel eki-791612FF
                                 --ramdisk eri-CFBE1450
```

- **Check instances:**

```
# ec2-describe-instances
RESERVATION   r-3DDE07D9    admin     default
INSTANCE      i-4901084F    emi-1DE4116D    0.0.0.0    141.52.166.160
running0m1.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
running0m1.small   2009-05-13T13:50:10+0000   eki-791612FF   eri-CFBE1450
```

- **Terminate instances:**

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

AppScale

http://appscale.cs.ucsb.edu

- **Open-source implementation of the Google AppEngine Cloud computing interface from UC Santa Barbara**
- **AppEngine allows to run web applications in the Google infrastructure**
- **AppScale executes automatically and transparently over Cloud infrastructures such as Eucalyptus, the open-source implementation of the Amazon Web Services interfaces**
- **AppScale provides a Platform-as-a-Service (PaaS) Cloud infrastructure that allows users to deploy, test, debug, measure, and monitor Google AppEngine applications prior to deployment on Google's proprietary resources**

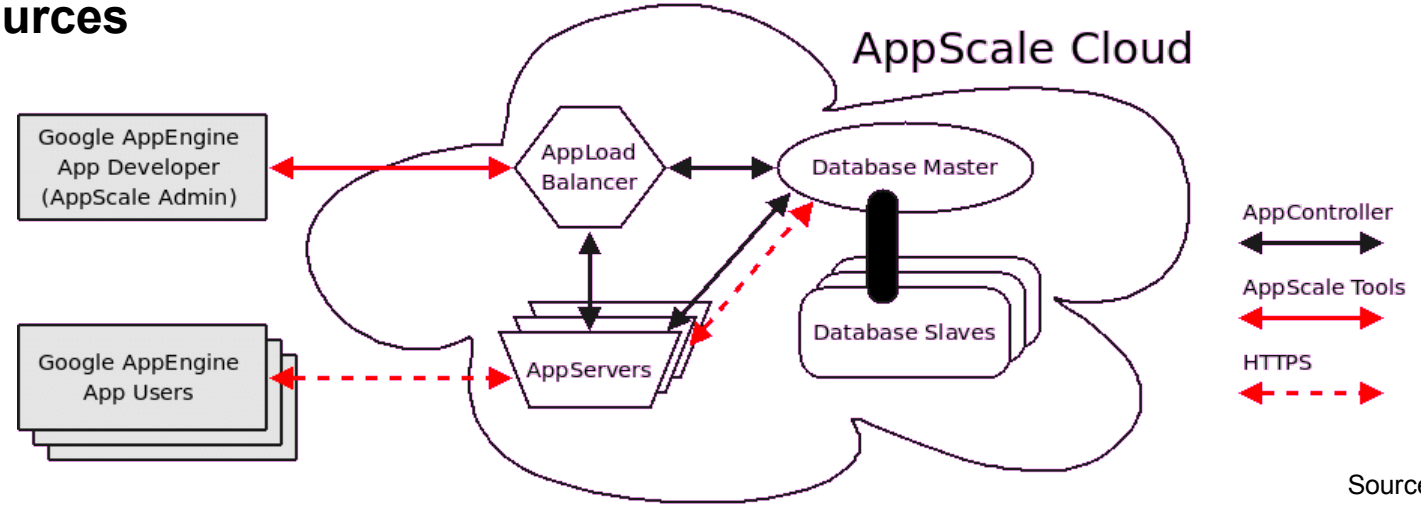

Source: Navraj Chohan

Summary

- Cloud computing is the next big thing
  - Flexible and elastic resource provisioning
  - Economy of scale makes it attractive
  - Move from manufacture towards industrialization of IT
  - (Everything as a Service)
- OpenCirrus™ offers interesting R&D opportunities
  - Cloud systems and application development
  - Accepting research proposals soon

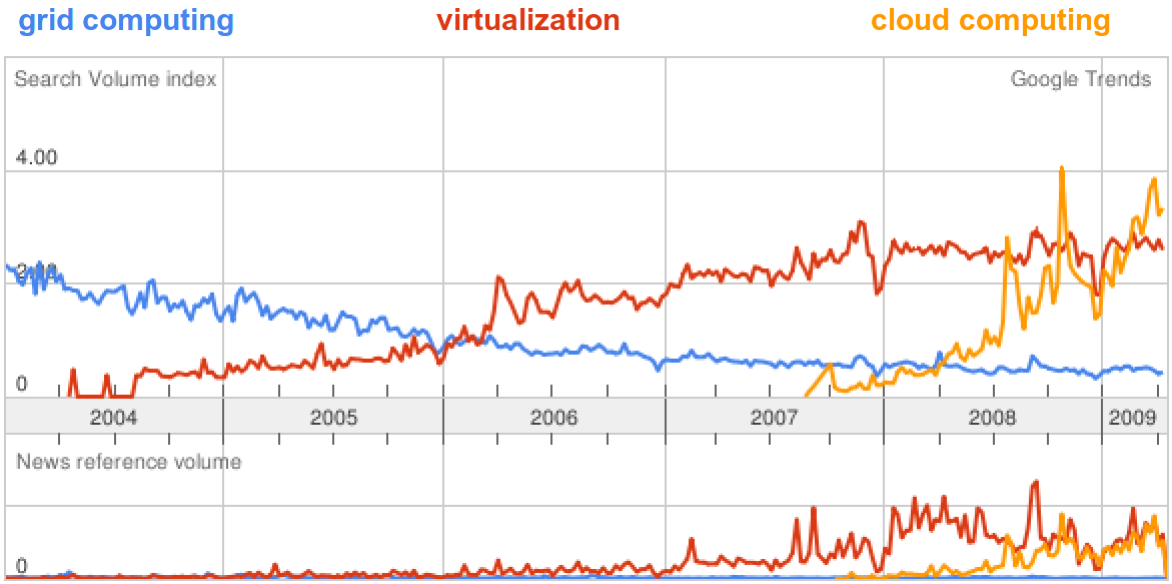

Thank you for your attention