

The OpenCirrus™ Project Cloud Testbed at SCC/KIT


Marcel Kunze
Steinbuch Centre for Computing (SCC)
Karlsruhe Institute of Technology (KIT)
Germany



Forschungszentrum Karlsruhe
in der Helmholtz-Gemeinschaft



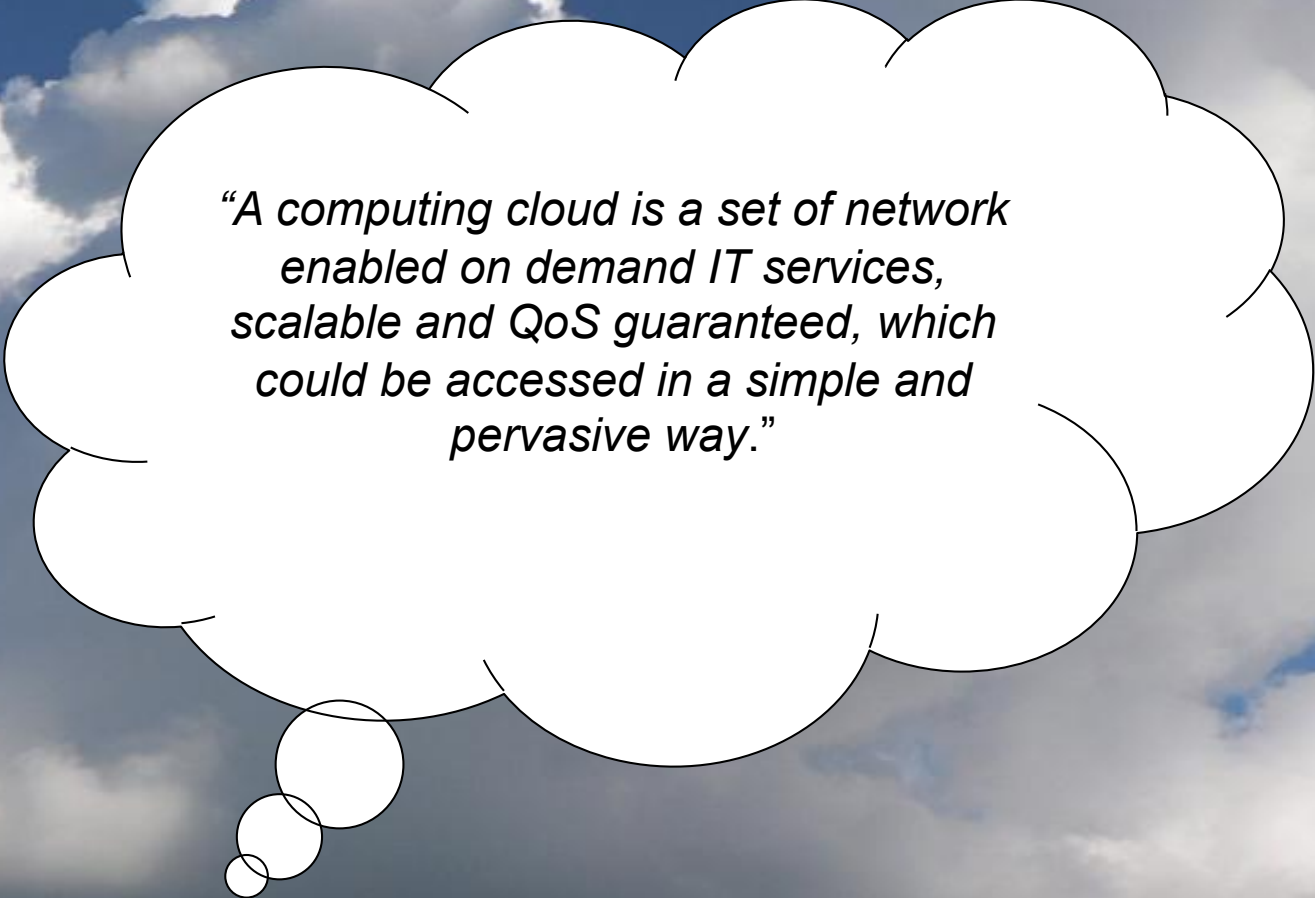
Universität Karlsruhe (TH)
Forschungsuniversität • gegründet 1825



Agenda

- What is cloud computing ?
- OpenCirrus™ project
- HPC and big data
- Summary

Cloud Computing: A possible Definition




“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.”

Commercial Cloud Offerings (Small Excerpt)



Amazon Elastic Compute Cloud (Amazon EC2) - Beta



Cloud Computing Overview
Cloudware - Cloud Computing Without Compromise

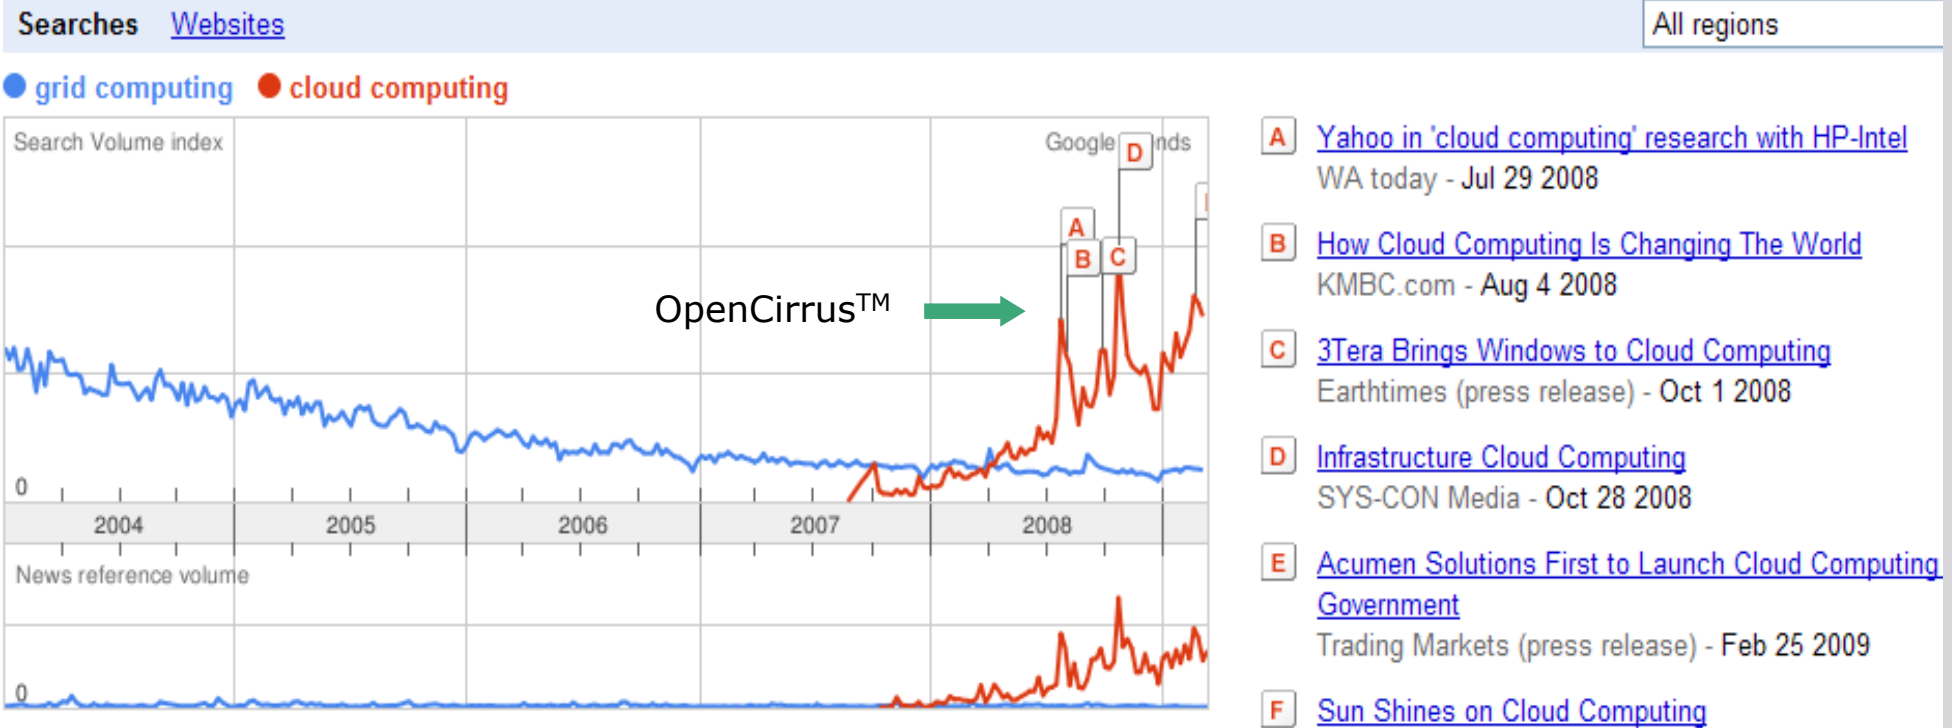



- **Problem: Commercial offerings are proprietary and usually not open for cloud systems research and development**

- **Simple, transparent, controllable cloud computing infrastructure**
 - What types of interfaces are appropriate for clouds?
 - How should cloud networks be constructed/managed?
 - How are security concerns addressed in “the cloud”?
 - How are various workloads most efficiently transferred?
 - What types of applications can run in clouds?
 - What types of service level agreements are appropriate/possible?

- **Research requirements**
 - Perform experiments also on a low system level
 - Flexible cloud computing framework
 - Compare different methodologies and implementations

Cloud Computing: A new Hype following Grid



■ Cloud computing R&D: OpenCirrus™ project

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.

OpenCirrus™ Cloud Computing Research Testbed

<http://opencirrus.org>

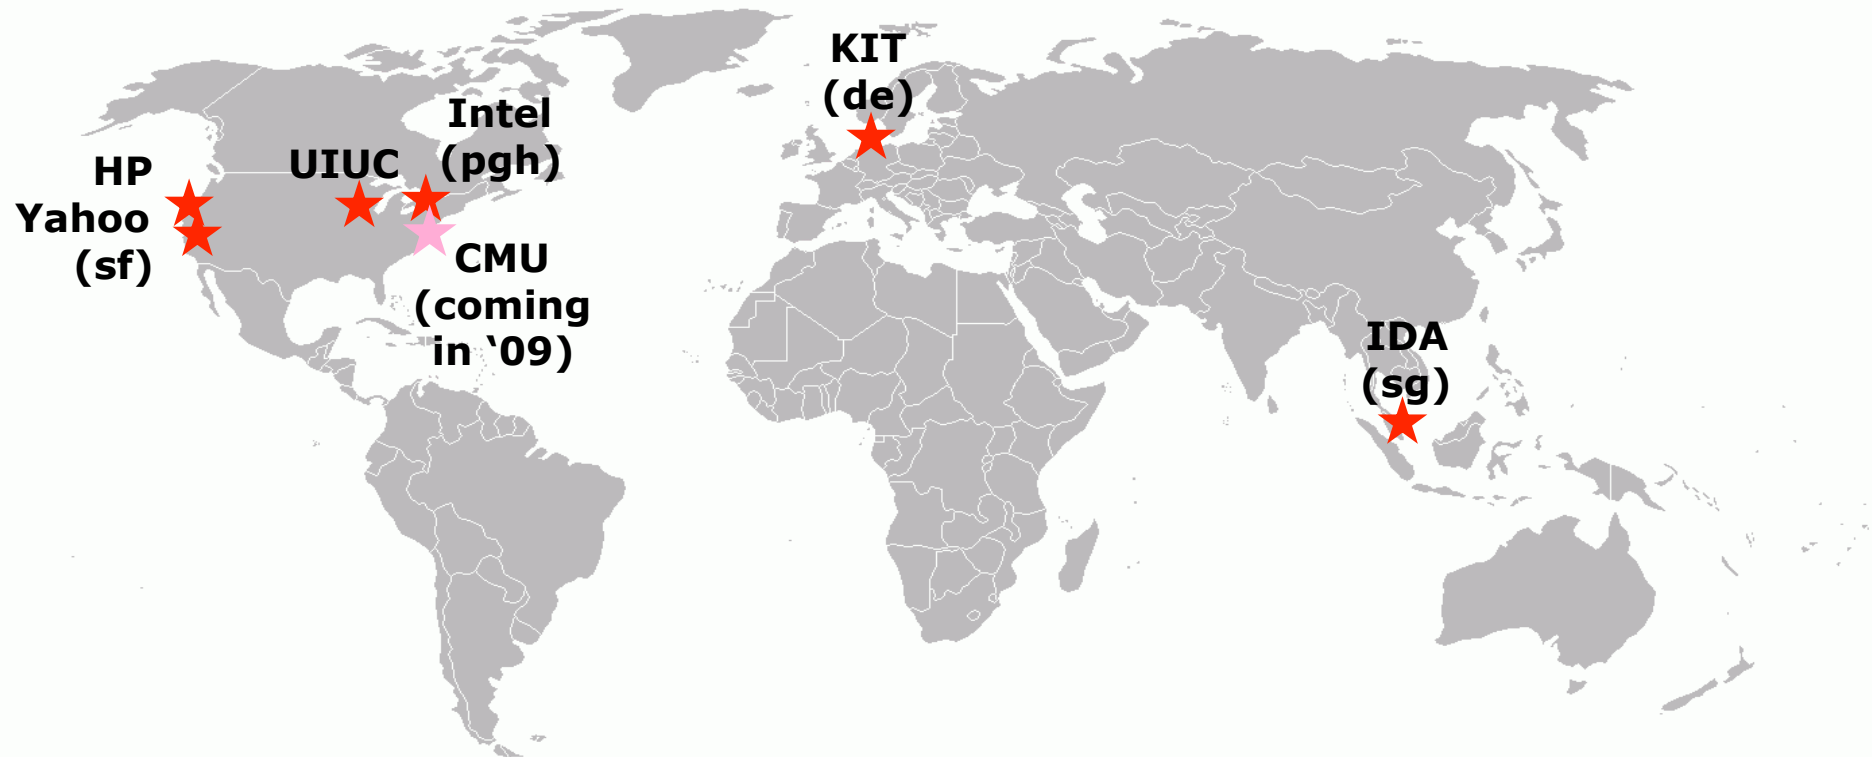


- An open, internet-scale global testbed for cloud computing research
 - Data center management & cloud services
 - Systems level research
 - Application level research
- Structure: a loose federation
 - Sponsors: HP Labs, Intel Research, Yahoo!
 - Partners: UIUC, Singapore IDA, KIT, NSF
 - Members: System and application development
- Great opportunity for cloud R&D

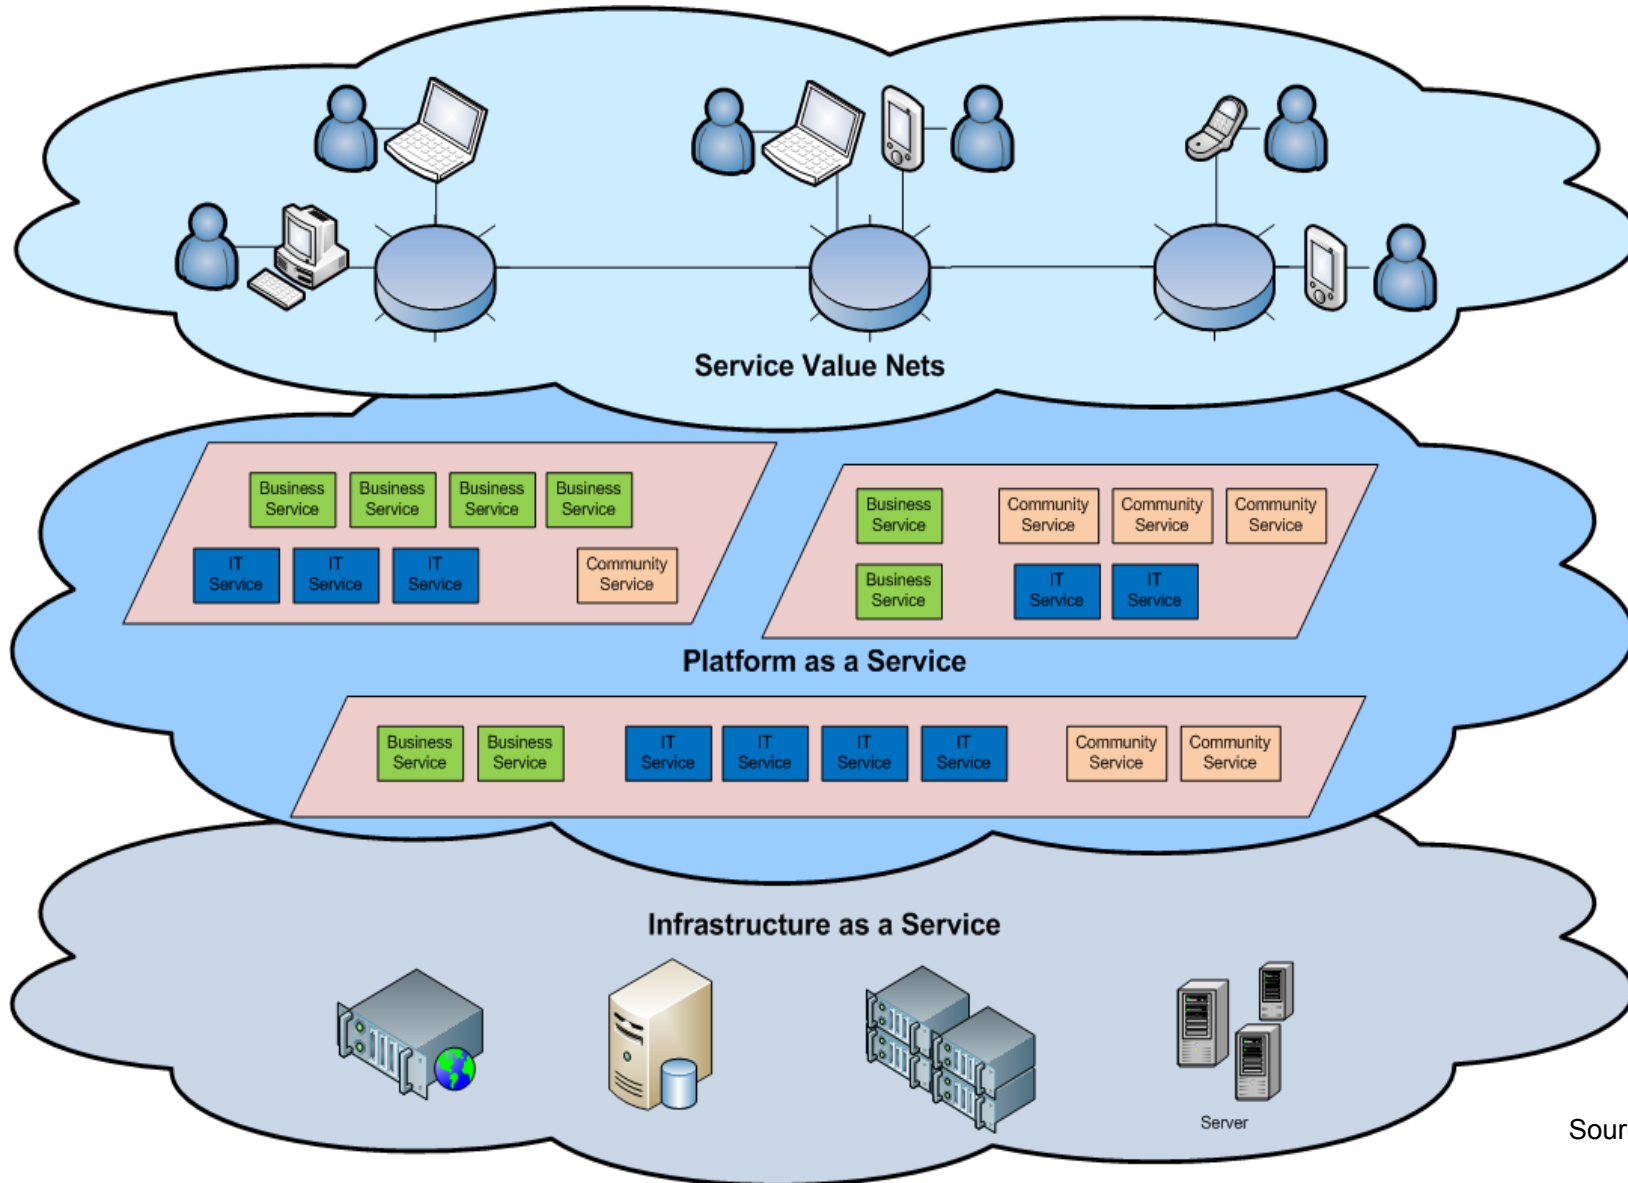


Where are the OpenCirrus™ sites?

- Six sites initially:
 - Sites distributed world-wide: HP Research, Yahoo!, UIUC, Intel Research Pittsburgh, KIT, Singapore IDA
 - 1000-4000 processor cores per site
 - SCC will host $o(3300)$ Nehalem T2 cores

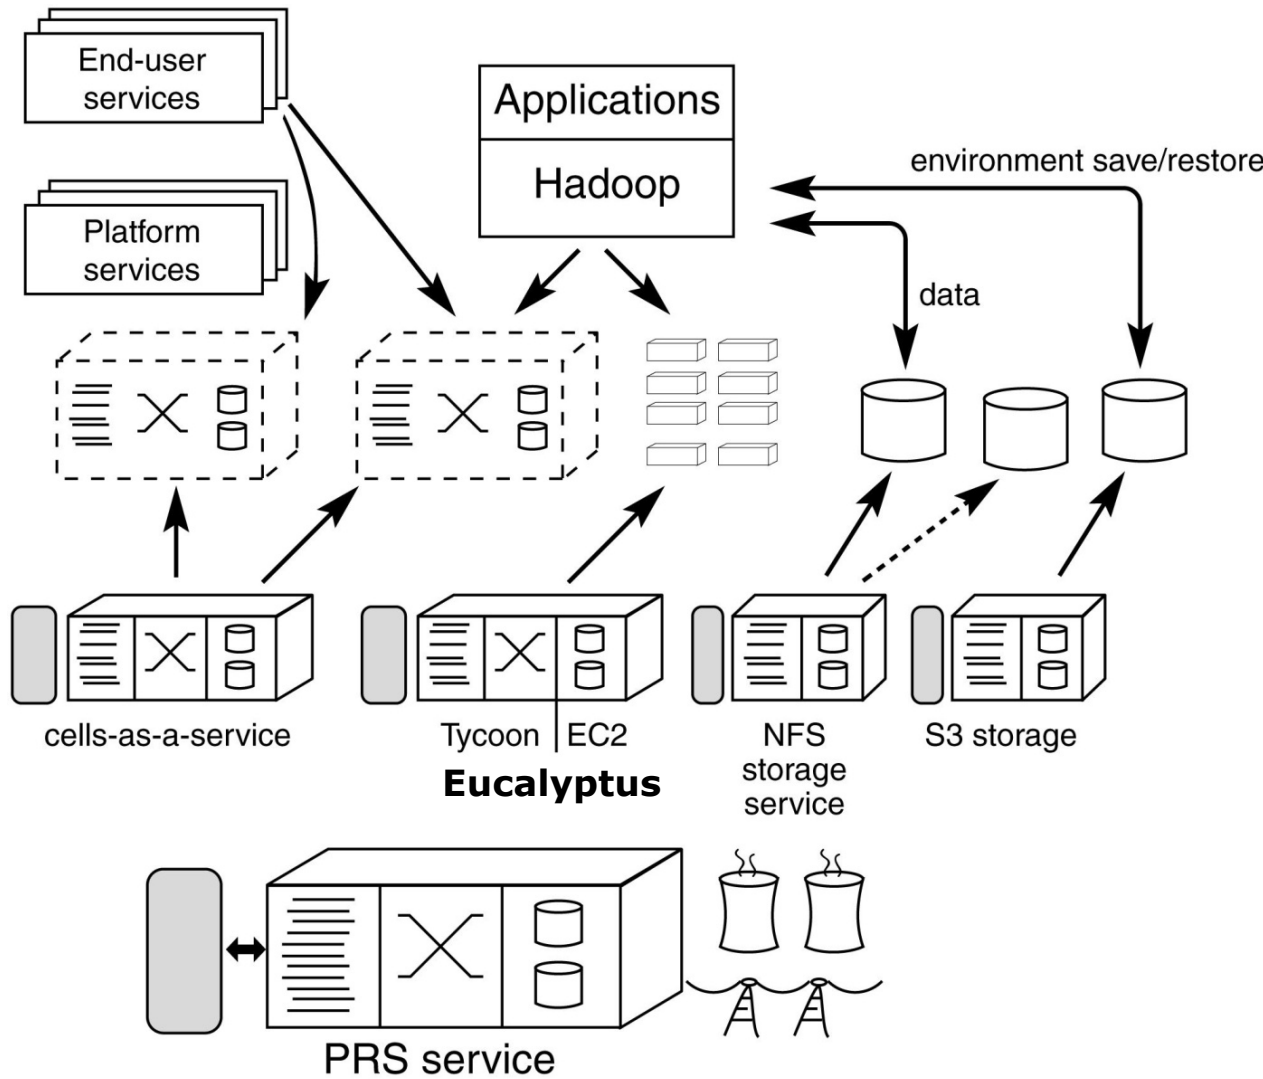


Cloud Architecture



Source: S.Tai

OpenCirrus™ Blueprint



Cloud application services

Virtual Resource Sets

Cloud infrastructure services

IT infrastructure layer
(Physical Resource Sets)

Physical Resource Sets (PRS)

■ PRS service goals

- Provide mini-datacenters to researchers
- Isolate experiments from each other
- Stable base for other research

■ PRS service approach

- Allocate sets of physical co-located nodes, isolated inside VLANs.
- Leverage existing software (e.g. Utah Emulab, HP OpsWare)
- Start simple, add features as we go
- Base to implement virtual resource sets

■ Hardware as a Service (HaaS)

Virtual Resource Sets (VRS)

- Basic idea: Abstract from physical resource by introduction of a virtualization layer
- Concept applies to all IT aspects: CPU, storage, networks and applications, ...
- Main advantages
 - Implement IT services **exactly** fitting customer's varying need
 - Deploy IT services on demand
 - Automated resource management
 - Easily guarantee service levels
 - Live migration of services
 - Reduce both: CapEx and OpEx
- Infrastructure as a Service (IaaS)
 - Implement Compute and Storage services
 - De-facto standard: Amazon Web Services interface

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 basic implementation of Amazon S3 interface
- Potential to interact with the same tools, known to work with Amazon EC2 and S3
- Eucalyptus is an important step to establish an open Cloud computing infrastructure standard

Eucalyptus

<http://eucalyptus.cs.ucsb.edu>



Amazon EC2 and S3 Interface


Client-side API Translator

Cloud Controller

Database

Cluster Controller

Cluster Controller



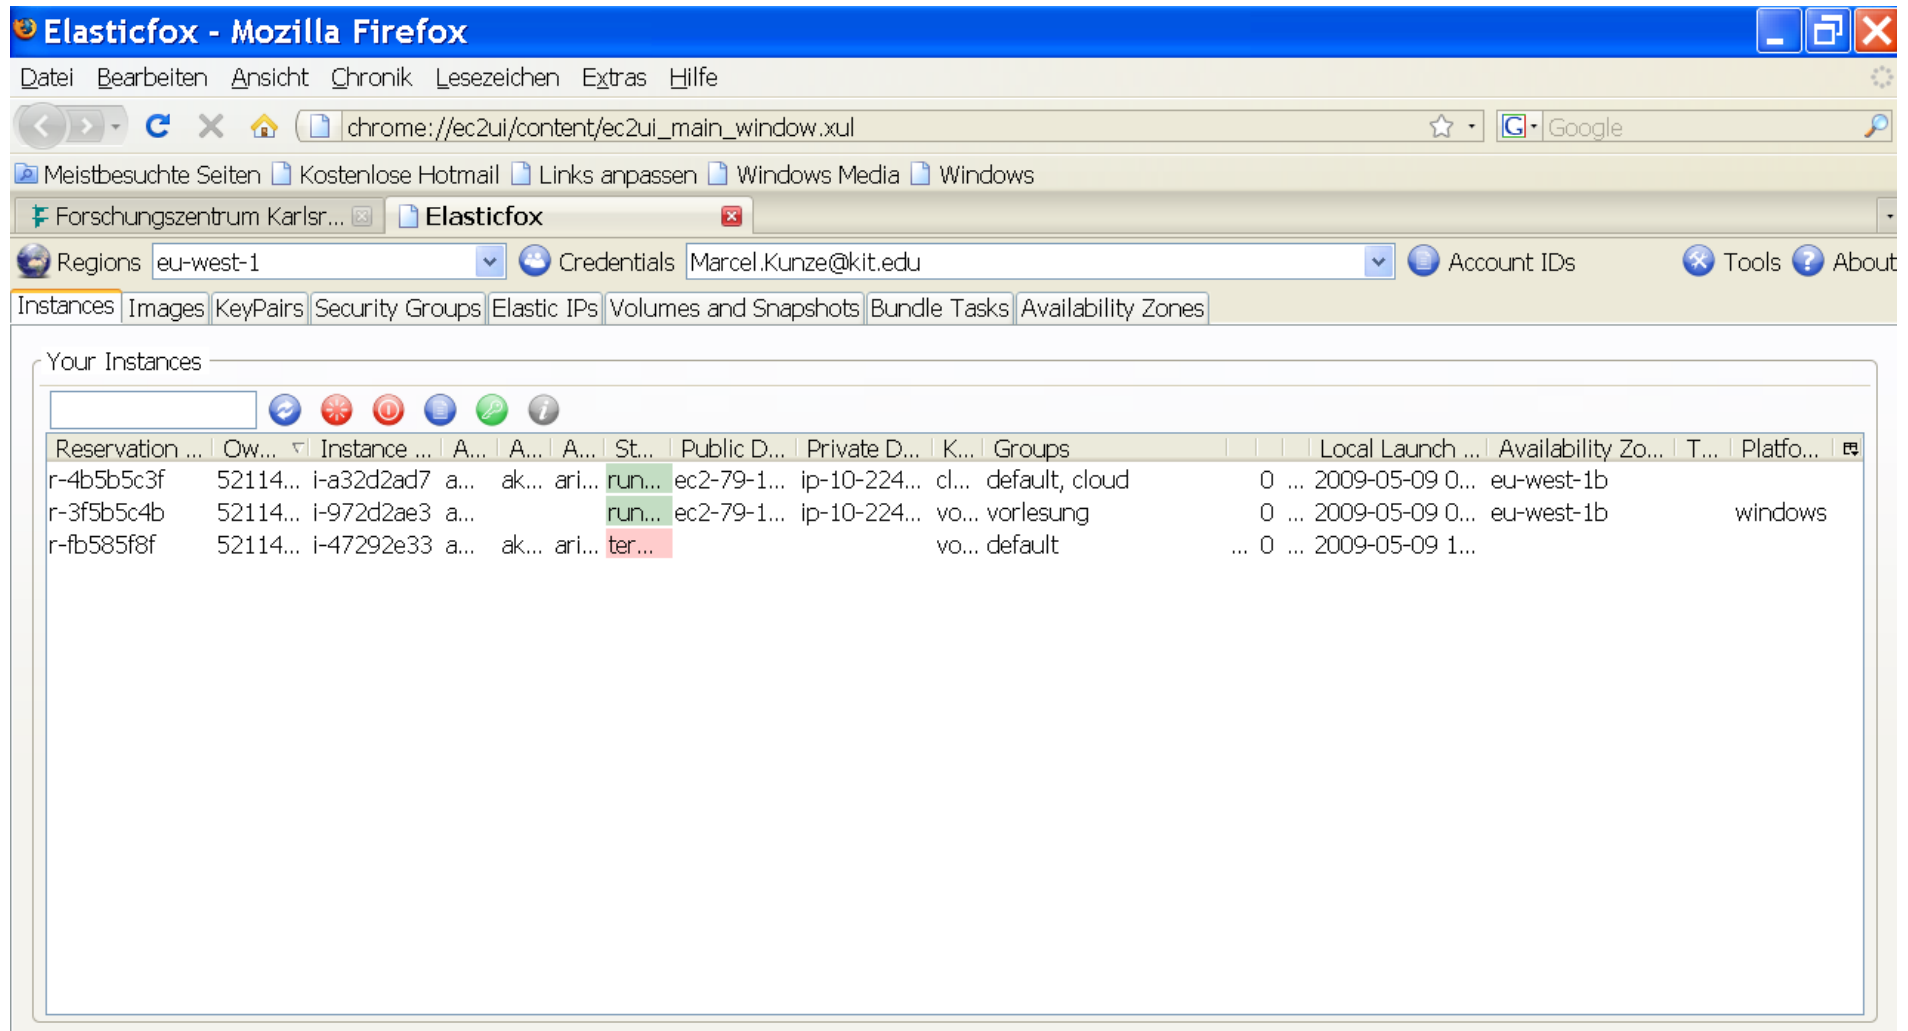
Schedules the distribution of virtual machines to the NC. Collects (free) resource information.

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

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

Source: R.Wolski

IaaS Self Service Portal with ElasticFox

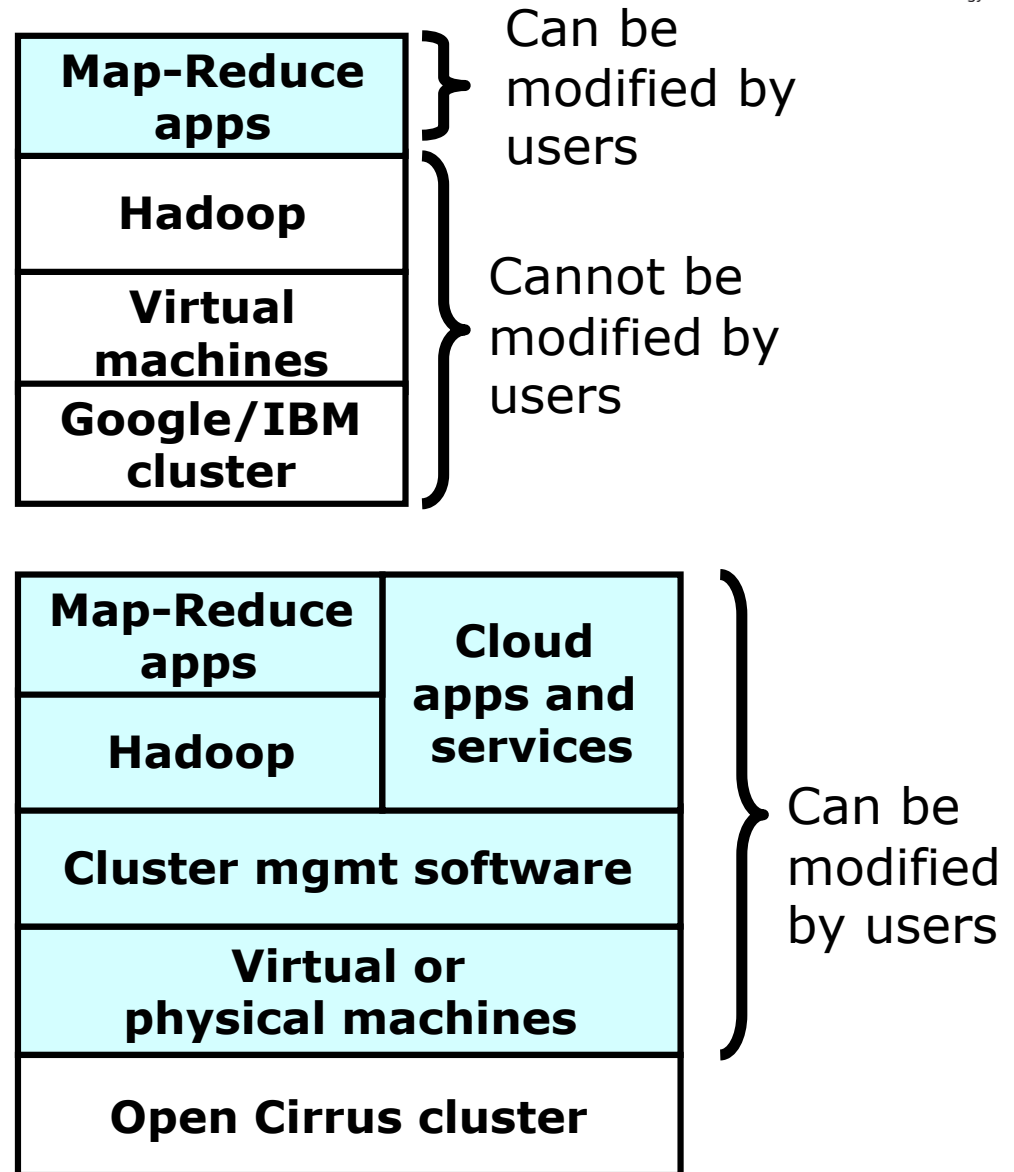


| Reservation ... | Ow... | Instance ... | A... | A... | A... | St... | Public D... | Private D... | K... | Groups | Local Launch ... | Availability Zo... | T... | Platfo... |
|-----------------|----------|--------------|------|-------|--------|--------|-------------|--------------|-------|----------------|------------------|--------------------|------------|-----------|
| r-4b5b5c3f | 52114... | i-a32d2ad7 | a... | ak... | ari... | run... | ec2-79-1... | ip-10-224... | cl... | default, cloud | 0 ... | 2009-05-09 0... | eu-west-1b | |
| r-3f5b5c4b | 52114... | i-972d2ae3 | a... | | | run... | ec2-79-1... | ip-10-224... | vo... | vorlesung | 0 ... | 2009-05-09 0... | eu-west-1b | windows |
| r-fb585f8f | 52114... | i-47292e33 | a... | ak... | ari... | ter... | | | vo... | default | ... 0 ... | 2009-05-09 1... | | |

■ Plan: Adapt ElasticFox FireFox Plugin for use with Eucalyptus

How is OpenCirrus™ different from other testbeds?

- OpenCirrus™ supports both system- and app-level research
 - n/a at Google/IBM and EC2/S3
 - OpenCirrus™ researchers will have complete access to the underlying hardware and software platform.
 - OpenCirrus™ allows Intel platform features that support cloud computing (e.g. DCMI, NM) to be exposed, and exploited.



How do users get access to OpenCirrus™ sites?

- Project PIs apply to each site separately.
- Contact names, email addresses, and web links for applications to each site will be available on the OpenCirrus™ Web site
 - <http://opencirrus.org>
- Each OpenCirrus™ site decides which users and projects get access to its site.
- Planning to have a *global sign on* for all sites
 - Users will be able to login to each OpenCirrus™ site for which they are authorized using the same login and password.

What kinds of research projects are OpenCirrus™ sites looking for?

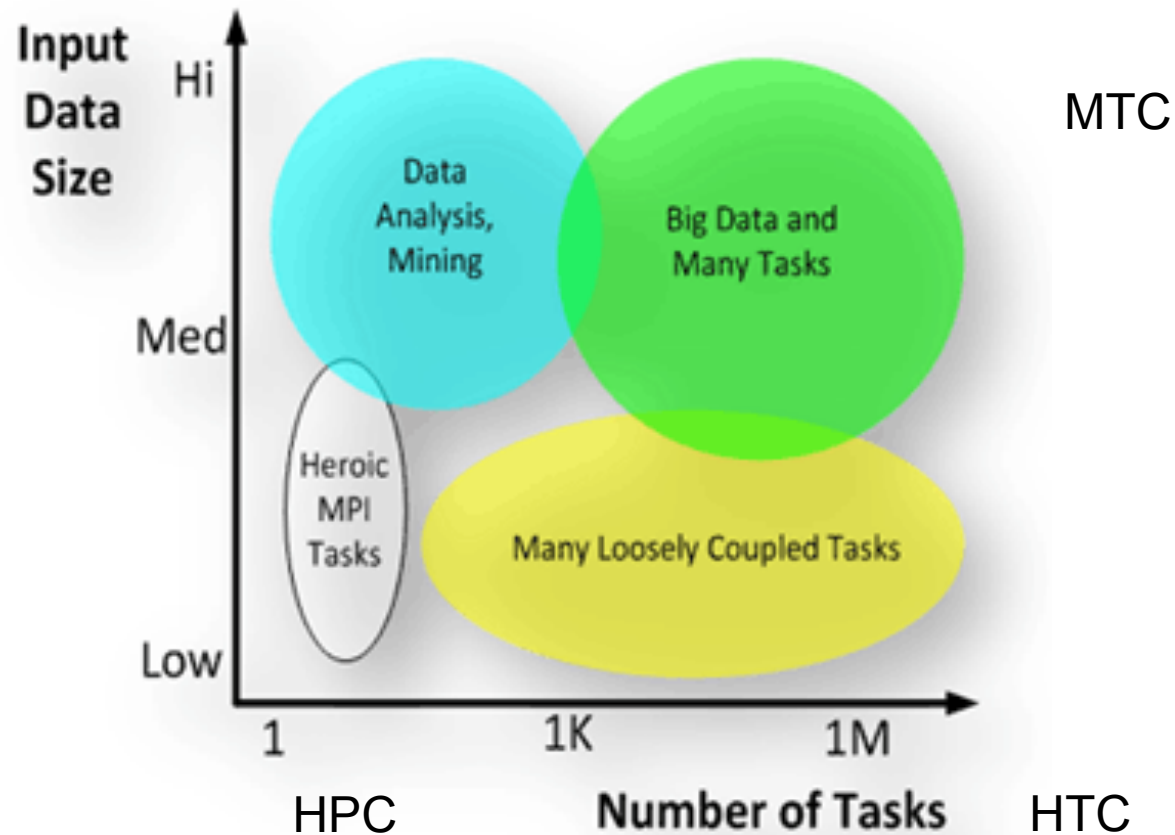
- Open Cirrus™ is seeking research in the following areas (different centers will weight these differently):
 - Datacenter federation
 - Datacenter management
 - Web services
 - Data-intensive applications and systems
 - Hadoop map-reduce applications

- The following kinds of projects are not of primary interest:
 - Traditional HPC application development.
 - Production applications that just need lots of cycles.
 - Closed source system development.

Potential Fields of Cloud System Development

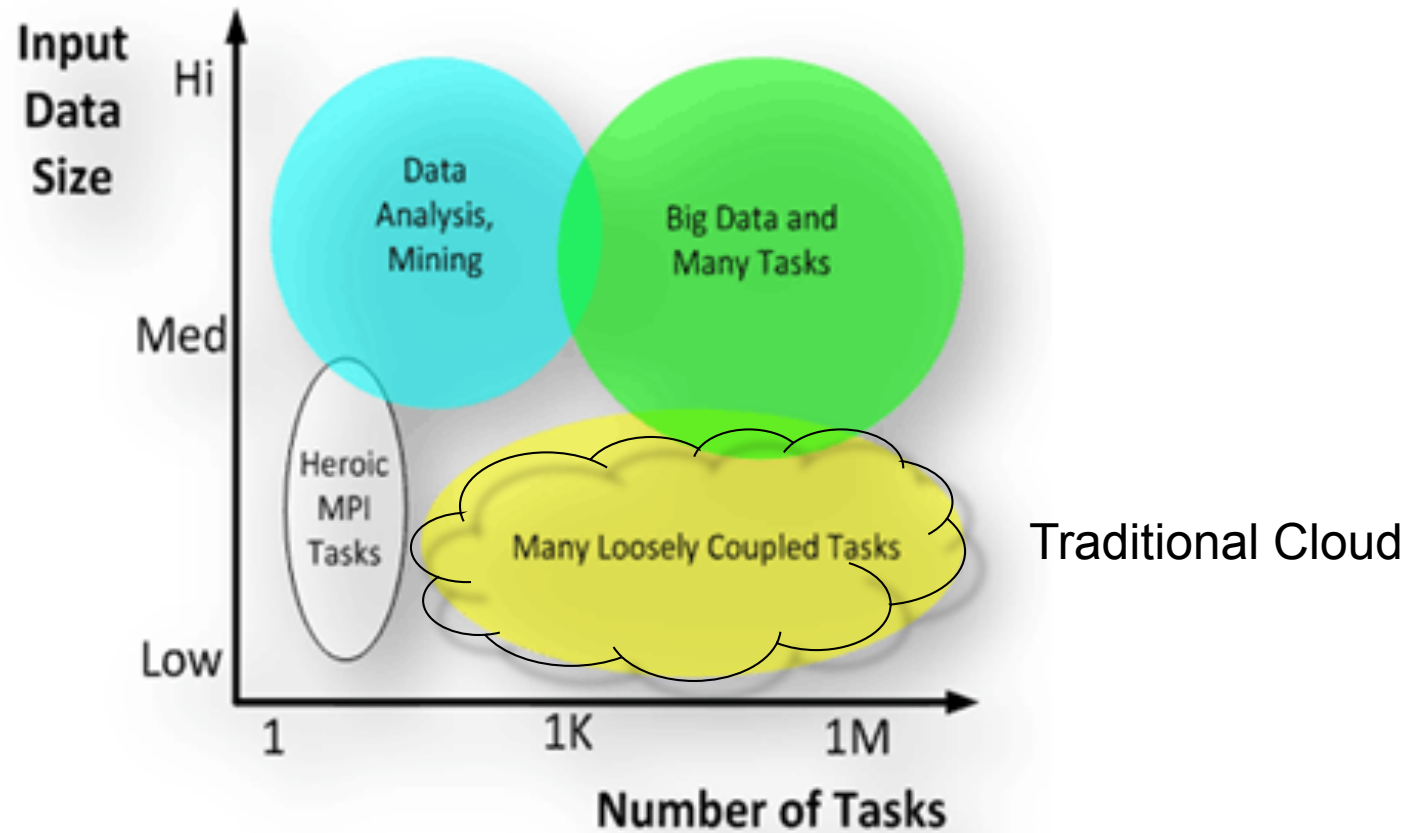
- New infrastructure services
 - HPCaaS: High Performance Computing as a Service
 - LSDFaaS: Large Scale Data Facility as a Service
 - GenomeDBaaS: Genome Database as a Service

HPC vs. HTC vs. MTC (Many Task Computing)




Source: I.Foster

The Cloud Space



Extension of the Cloud Space to all Areas



HPCaaS

- High Performance Computing as a Service
- Interesting Fields for R&D in Open Cirrus™
 - Flexible platform services for HPC customers
 - Development of MPI services for clouds
 - Development of scheduling services for clouds
 - Management of software licenses
 - Integration of Grid resources: Grid as a Service (GaaS)

LSDFaaS

- Large Scale Data Facility as a Service
- Actual projects at KIT in this field:
 - Data storage for LHC computing
 - Data storage for ITER (EUFORIA)
 - Project ANKA (synchrotron radiation source)
 - Activities in materials research
- Long-term data filing due to legal requirements
- Development of big data services

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 development
 - Cloud application development
 - Accepting research proposals soon
- OpenCirrus™ summit at HP Palo Alto on June 8/9

Karlsruhe Institute of Technology



Steinbuch Centre for Computing (SCC)

Thank you for your attention.

