

A VISION for Cloud Computing

Per Brand, SICS

Outline

- Cloud computing
 - What it is
 - Who benefits (most)
 - Advantages and drawbacks
 - Research Issues
- VISION – EU IP project on clouds
 - Project Goals
 - SICS role

What is Cloud Computing

- Computing power and specific applications will be sold using a utility business model (e.g., water, electricity) – John McCarthy 1961
- A cloud is an **elastic** execution environment of resources ... providing a **metered service** at **multiple granularities** for a **specified QoS** level
EU Expert Group Report Version 1.0, Jan. 2010.

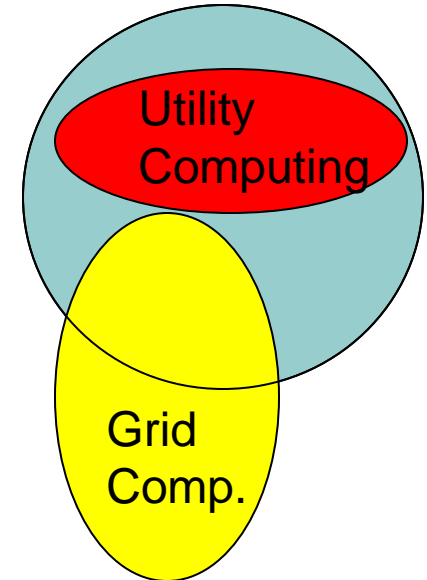
From 1961 to mid 1990s

- The 61 vision died and was largely forgotten by the 70's – impossible technically to realize
- Concentrating on the most basic utility – communications infrastructure (TCP, networking tech)
 - Computers & storage individually owned & managed
 - Even with this basic utility a lot could be done – e.g. WWW


From mid 1990s to 2010

- In the late 1990's idea resurfaced as **utility computing**
- 1990's Grid Computing – federating resources
- 2000's increasing focus on data-intensive applications
- Google and Amazon started the cloud revolution almost by accident because of their internal needs, but now it has its own momentum
- Important technologies
 - Virtualization (e.g. XEN 2003)
 - Other advances from autonomic computing & distributed systems/algorithms

Cloud Computing today




Origins of the Term - Cloud



Previously:
Communications
Infrastructure

Origins of the Term - Cloud




Previously:
Communications
Infrastructure

Clouds are
Cloudy:

Cannot see
internal details

Don't need to:
can use without
knowledge of how
or where

Origins of the Term - Cloud




Previously:
Communications
Infrastructure



Clouds are
Cloudy:


Cannot see
internal details

Don't need to:
can use without
knowledge of how
or where




Now:
Also data, computation

Some important properties of clouds


- Advantages
 - Elasticity !!
 - Scalability
 - Pay-per-use, i.e. OPEX not CAPEX
 - Economy of scale
 - No overprovisioning
 - Lower energy consumption
 - Drawbacks
 - Lack of standards
 - Vendor lock-in
 - Security
 - Legal Issues
 - Cost of data shipping
- Probably Solvable
- 

Four patterns where clouds are most beneficial

Resources




Rapid growth



Failed StartUp



No money
No credit

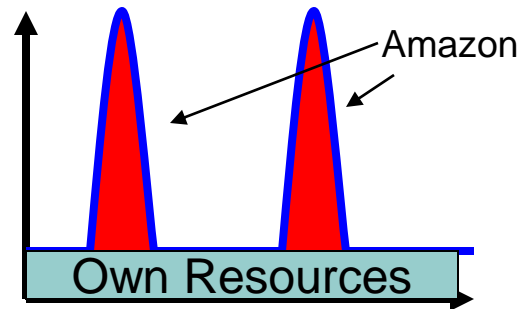


Variable Needs
Bursty Applications

There's burstiness and then there is burstiness

Some burstiness everywhere
(average cluster utilization
<20%)

Sometimes much more
peaks orders of magnitude
larger than average.



Example:

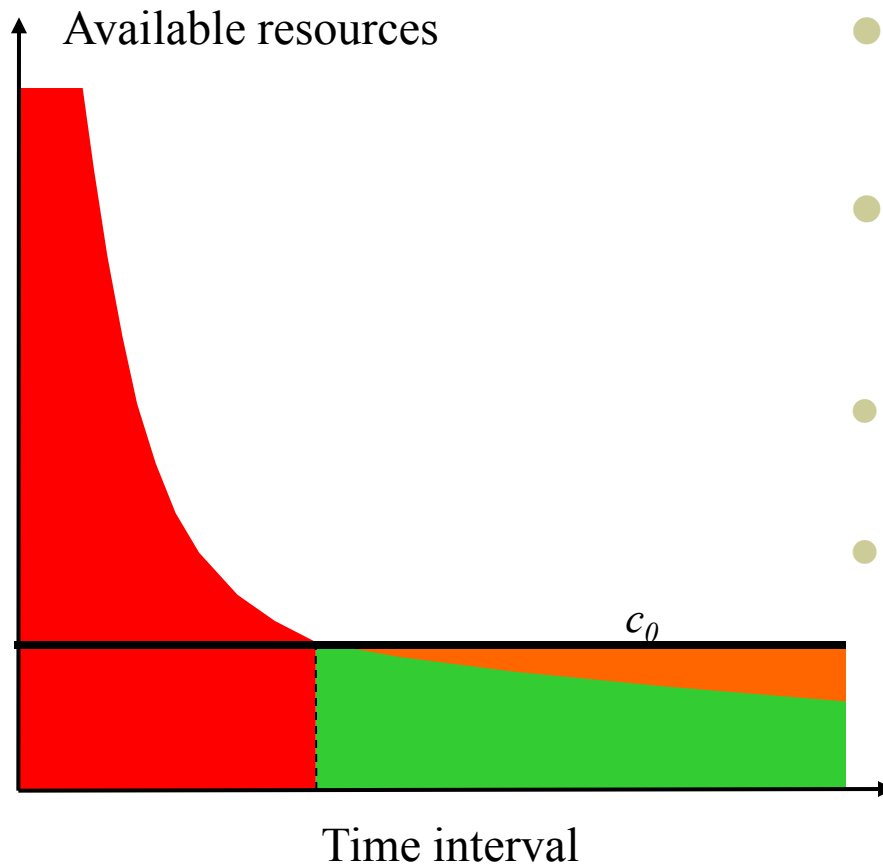
Chemistry – models that predict the properties of
chemicals from their structure

Irregular (e.g. once a month)

New chemical data

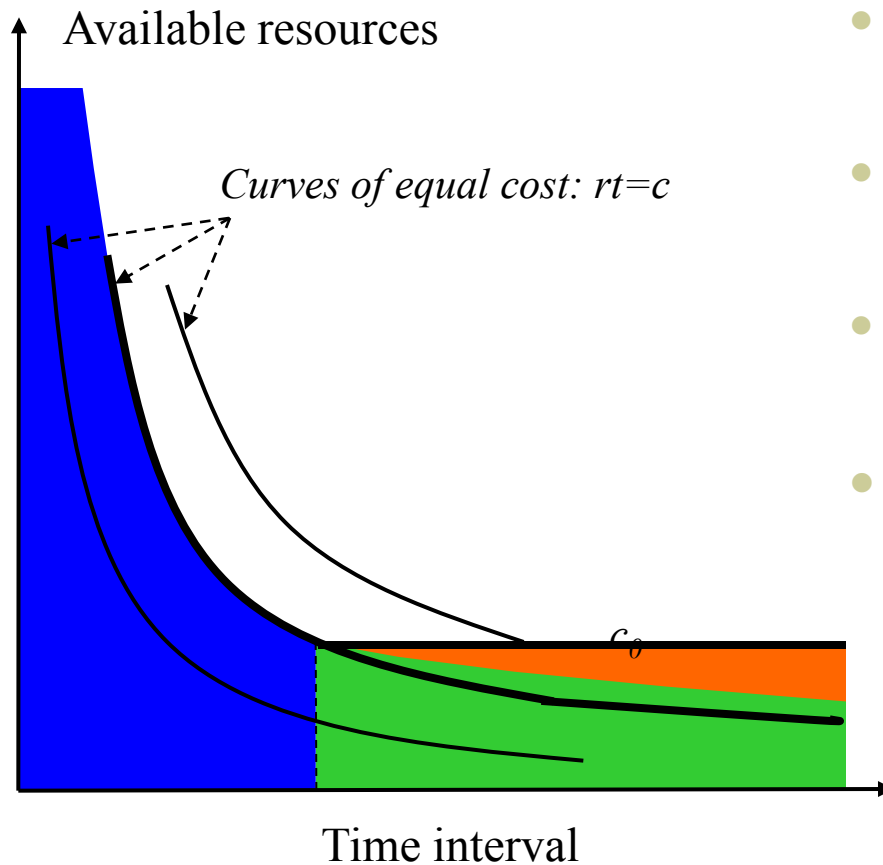
New improved models

Traditionally



- **Choose your cost level**
 - Amount of resources
- **Green=useful utilization**
- **Orange= overprovisioning**
- **Red=problems**
 - Break
 - Less useful work done
 -

The cloud revolution



- **Curves of equal cost with clouds**
- **Clouds allow you to move freely along the lines of equal cost**
- **To some degree – limits to elasticity even with clouds**
- **Questions**
 - Are there new kinds of applications lie on the extreme left ?
 - How far to the left can we go ?

VISION project

- Goal: Develop an architecture and a reference implementation of a cloud infrastructure for data-intensive applications
 - building on open standards
 - federating clouds
 - optimized for *data-intensive* applications.
- European IP project 2010-2013
 - Partners: IBM, Siemens, Deutsche Welle, RAI , Telefonica, FT, Telenor, Siemens, SAP, SICS, Engineering, Umeå, NTUA, iTricity, Messina, SNIA Europe
 - Use Cases: Media, Telco, Healthcare, Enterprise


Computational Storage in VISION

- Computation moved (near) to data
- Associated with (mobile) computations
 - Execution constraints (for safety and management)
 - Capabilities (for secure access)
 - Triggers (for activation) – computations may be long-lived, activated by data update.
- SICS will develop
 - The programming model
 - Expressiveness: Generalizing on Googles map/reduce and Microsoft Dryad
 - The runtime execution system
 - Moving and triggering computations
 - Enforcing execution constraints

A Swedish Cloud Computing WG

- In association with VISION, SICS will start a Cloud Computing Working Group
 - For Swedish industry & academia interested in cloud computing
 - Networking
 - Seminars, updates about developments in cloud computing
 - SMEs, startups welcome
 - Starting Sept 2010
 - We will contact those we know about
 - Otherwise contact Per Brand, perbrand@sics.se

Resources



Time