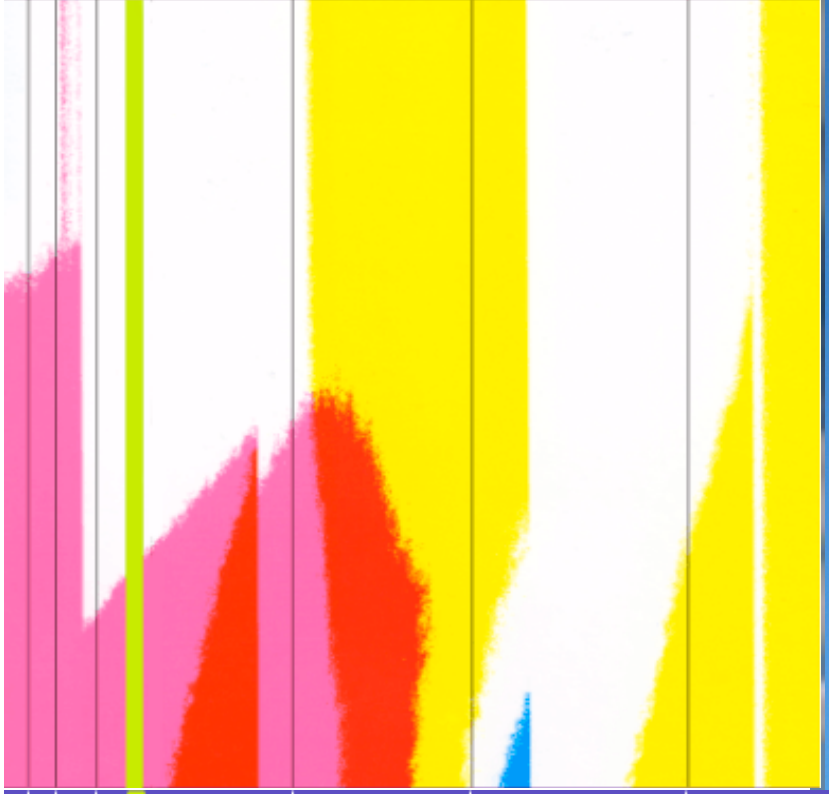

Cloud Computing

Gabriel Antoniu
Inria

INSTITUT NATIONAL
DE RECHERCHE
EN INFORMATIQUE
ET EN AUTOMATIQUE

INRIA

Computing as a Utility

first suggested by John McCarthy in 1961 !

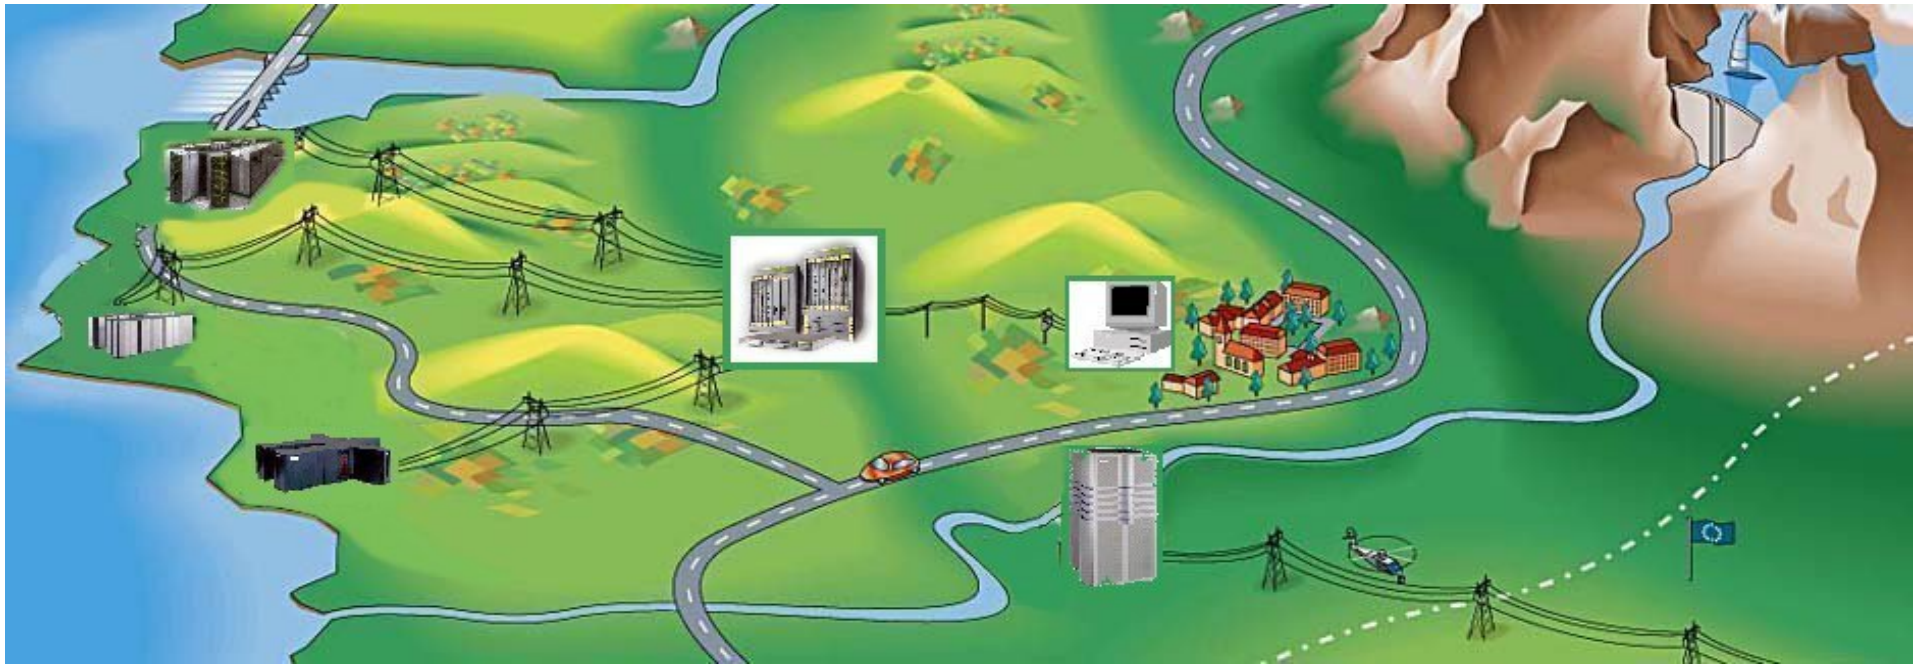

It is much cheaper to «rent» a computing infrastructure than building, operating and owning it !

Grid computing

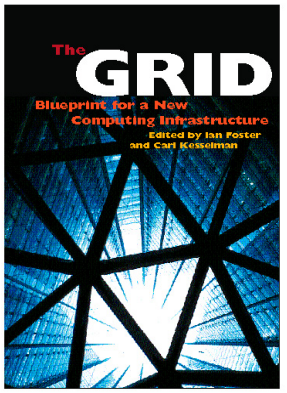

- What is Grid ?
  - *«A fully distributed, dynamically reconfigurable, scalable and autonomous infrastructure to provide location independent, pervasive, reliable, secure and efficient access to a coordinated set of services encapsulating and virtualizing resources (computing power, storage, instruments, data, etc.) in order to generate knowledge...» from the CoreGRID NoE*

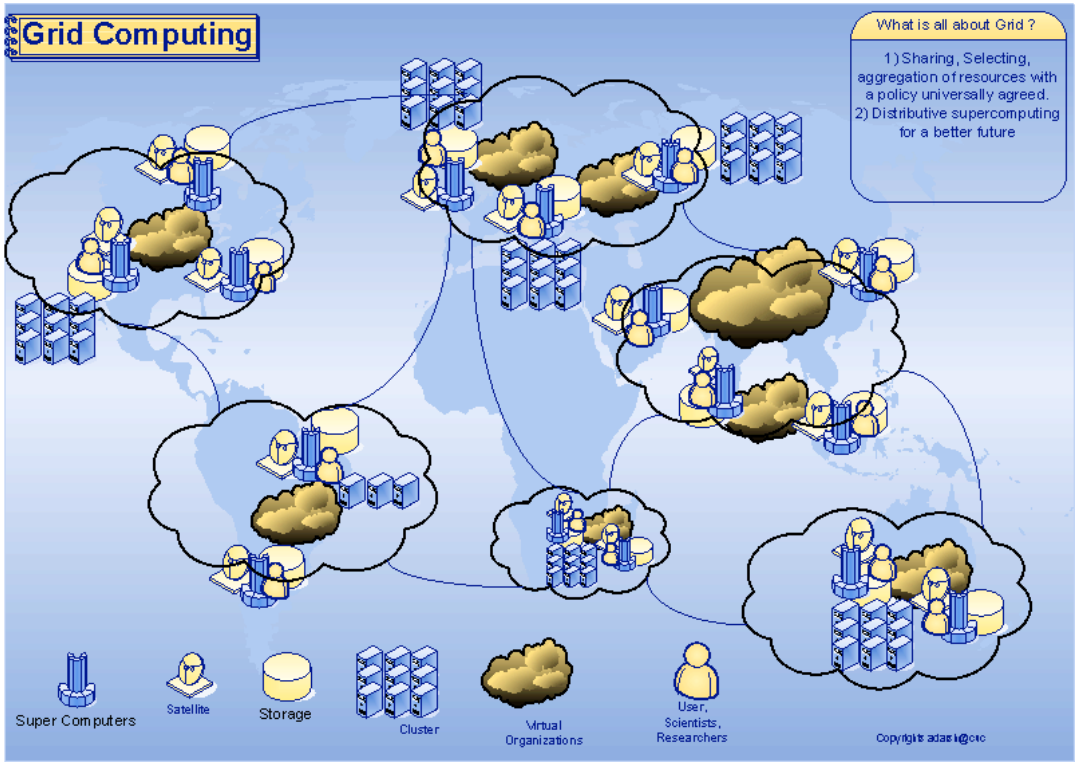

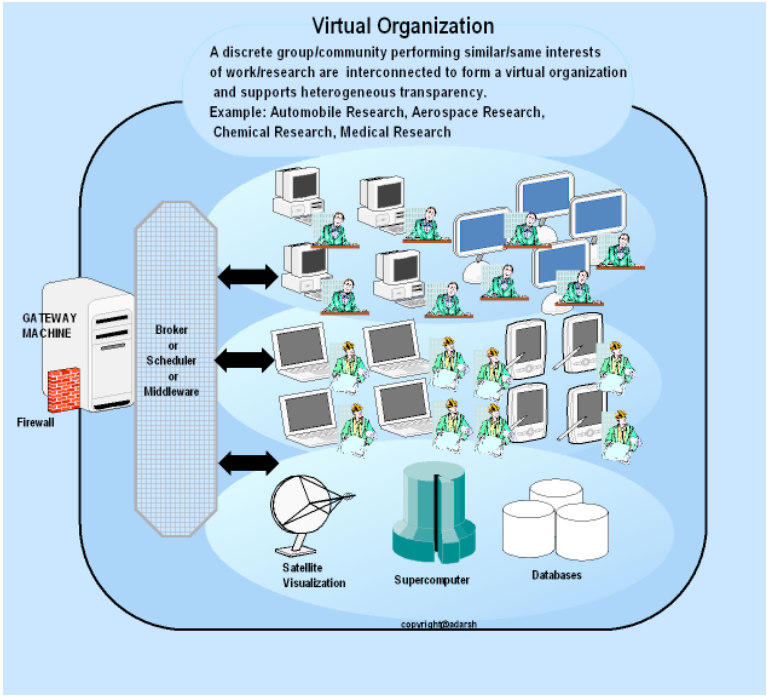

Cloud computing

- What is Cloud ?
  - *An emerging computing paradigm where applications, data and infrastructures are provided as a service that can be ubiquitously accessed from any connected devices over the internet.*

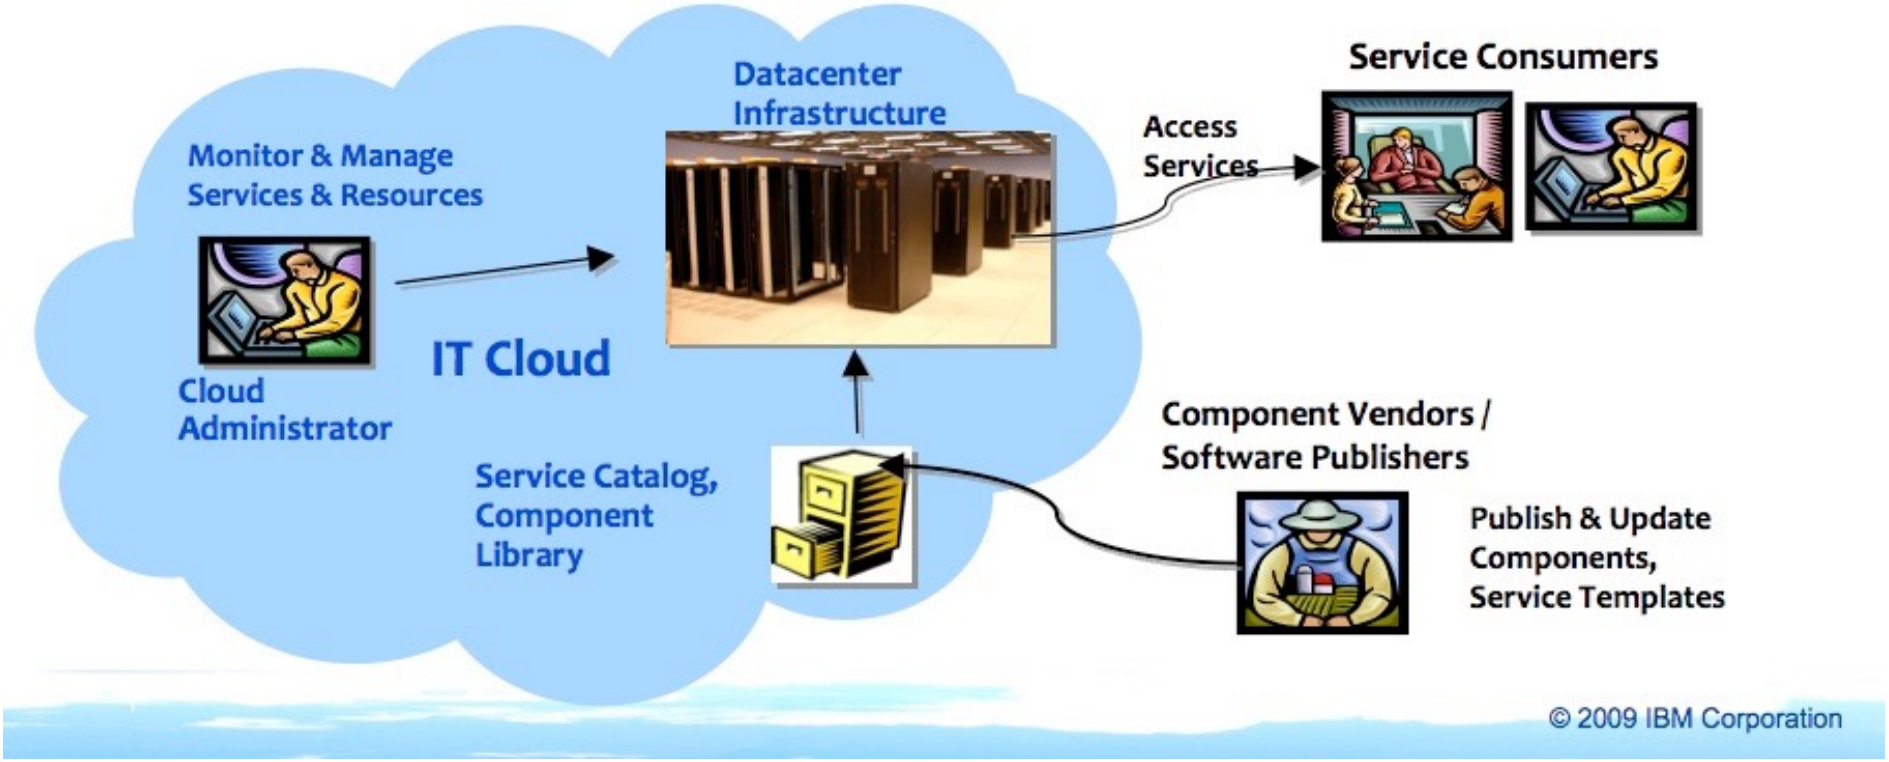

Cloud computing vs Grid Computing

- Distributed versus Centralized

- Resource provisioning
  - Batch scheduler / VMs management

What is behind Cloud

Google cluster 1997

- Datacenters as the reincarnation of the mainframe concept
- The end of the PC/Mac era ?
  - just a web browser is needed
  - «The network is the computer», «thin client», ...

Datacenters : easy to build !

- Based on the LEGO concept - a datacenter in shipping containers

- You do not even need a building, just gather these building blocks together on a parking lot and plug them to the Internet and to the power grid and that's it !
- Energy / Green-IT issues
  - In 25 years from now, Internet will consume the same quantity of energy than the humans today
  - Humans have to be ready to fight against computers to get access to the energy...

Datacenters : easy to build !

- If local laws matter... Google has a patent for this !

- Just set up offshore datacenter vessels out of territorial seas...

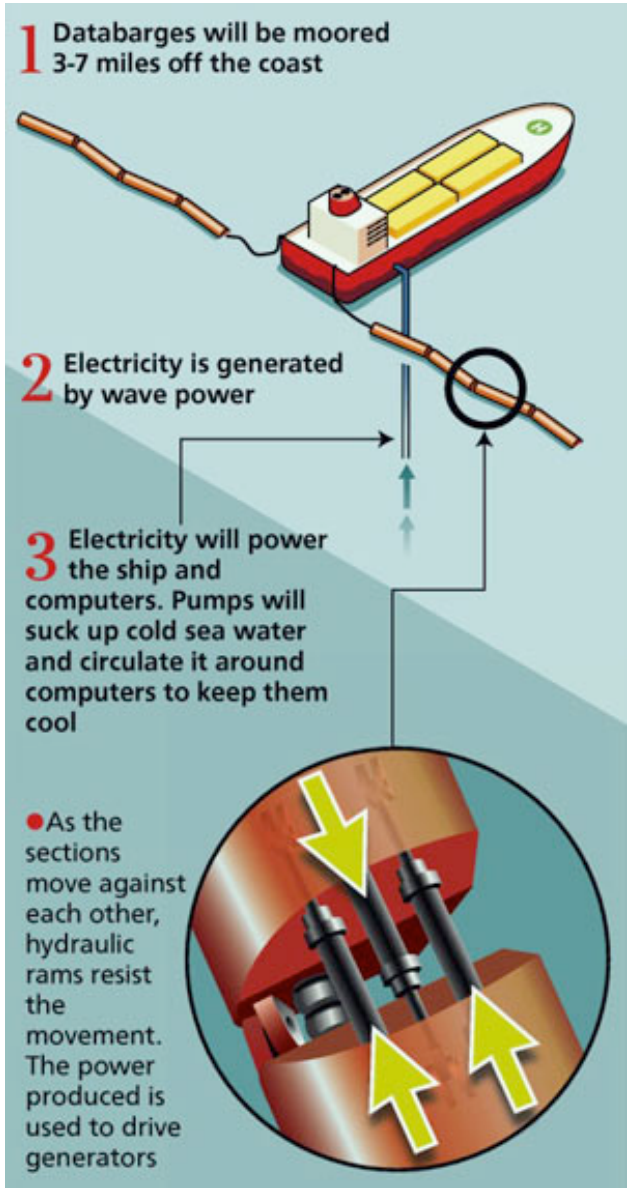

Image:

Why Cloud now and not before ?

- Internet !
  - Network performance has been improved dramatically the last 15 years
  - Nearly always connected to the Internet (anytime, anywhere)
- PC is not anymore the central device for personal computing
  - MP3, SmartPhone, Tablets, Set-top box, PCs, ...

  - How to get access to my personal data anywhere/anytime and from any devices ?
- Cost
  - Oversized systems to meet peak demand (both in the private and public sector)
  - Outsourcing (labor cost is much higher that computing cost)

Computing as a utility : a brief history

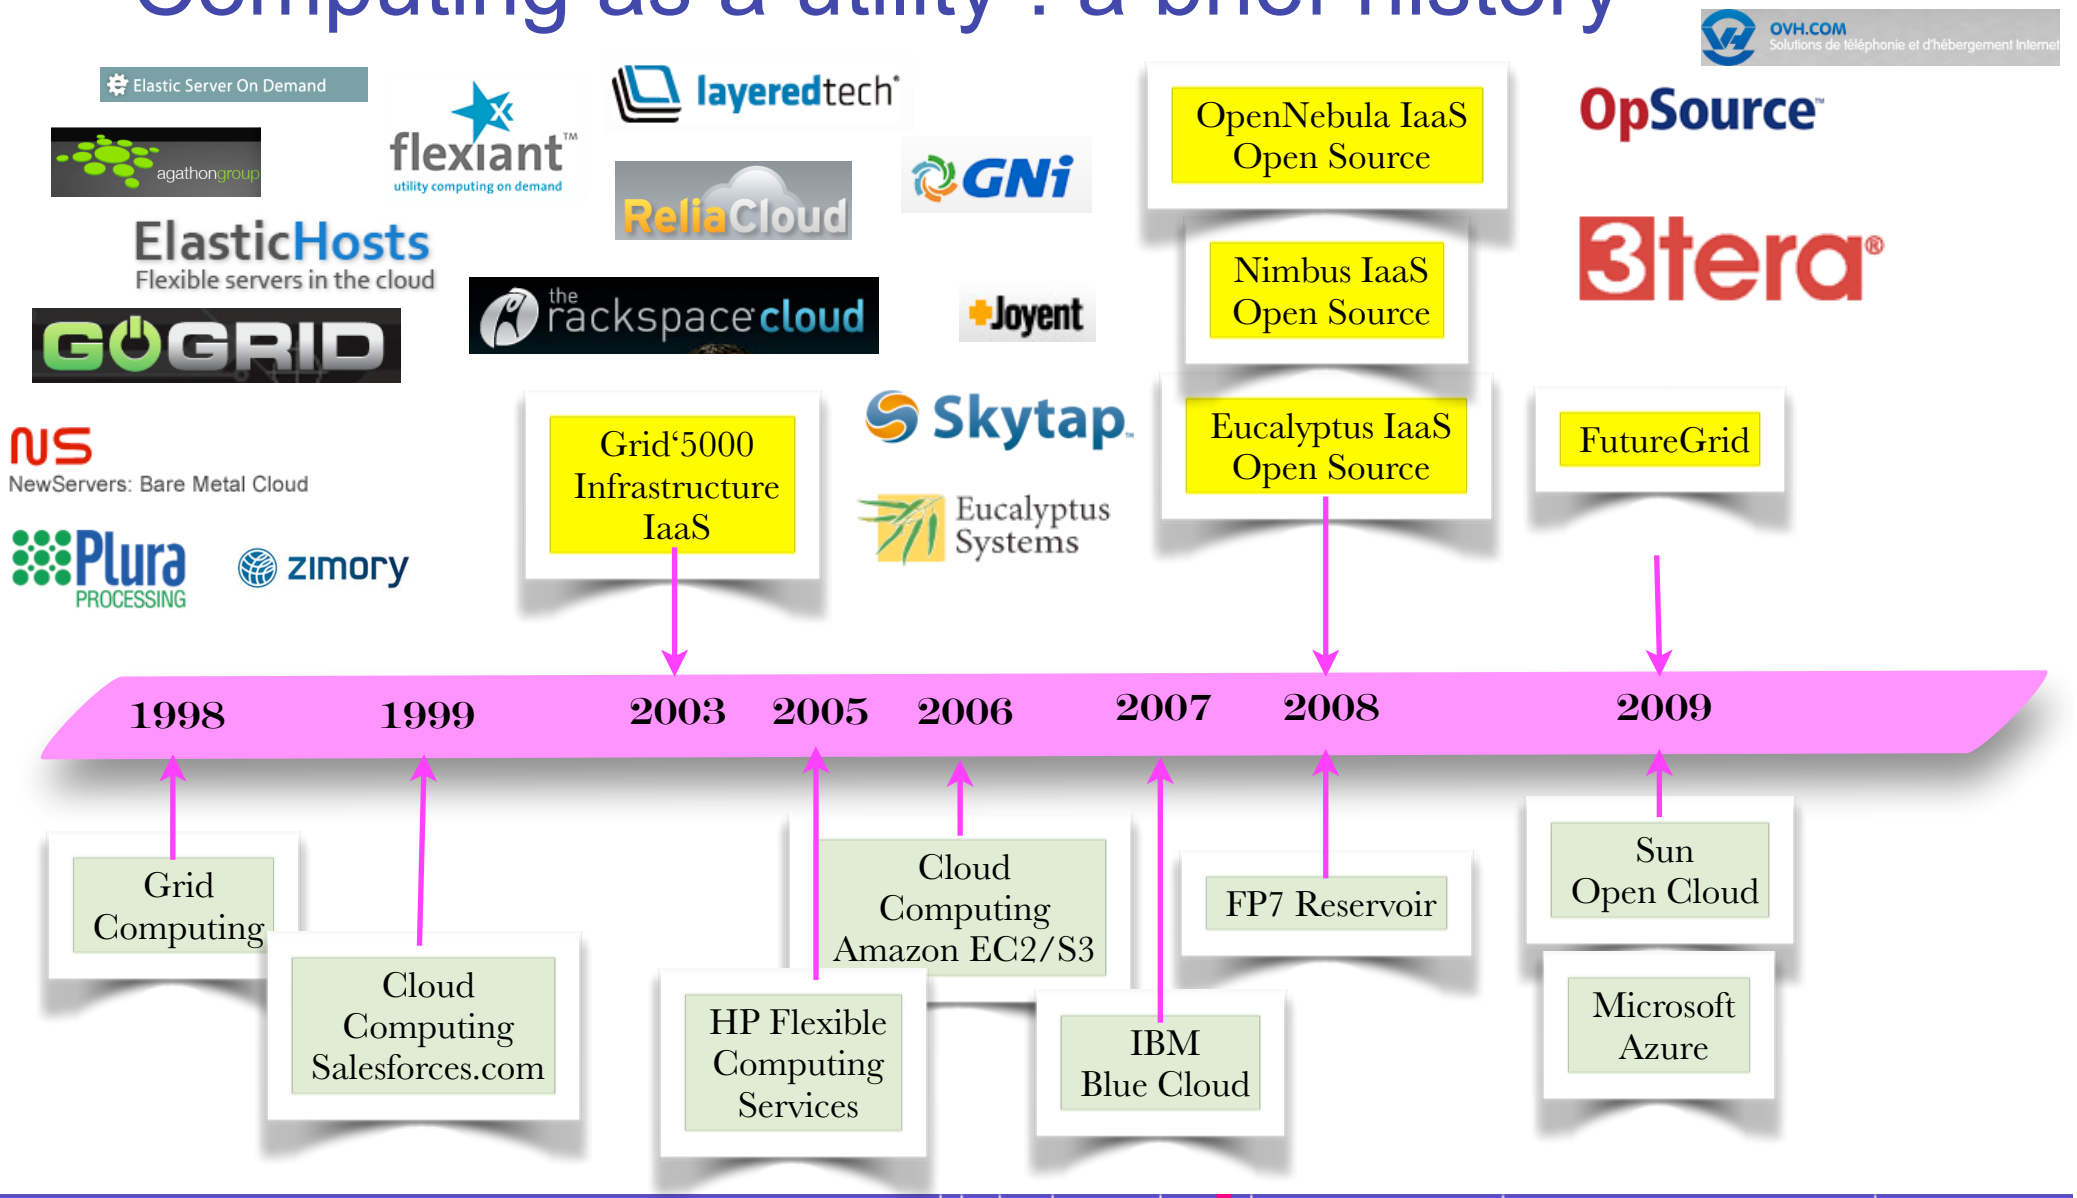

Cloud Acronyms

- PaaS - Platform/People as a Service
- SaaS - Software/Search as a Service
- IaaS - Infrastructure as a Service
- DaaS - Data as a Service
- CaaS - (composition/communication /composite) as a Service
- HaaS - Human as a Service ... just your shared agenda ;-)
- KaaS - Knowledge as a Service
- ...

- AaaS/XaaS - Anything as a Service or X to replace any letter...

Cloud: how to escape from the jungle

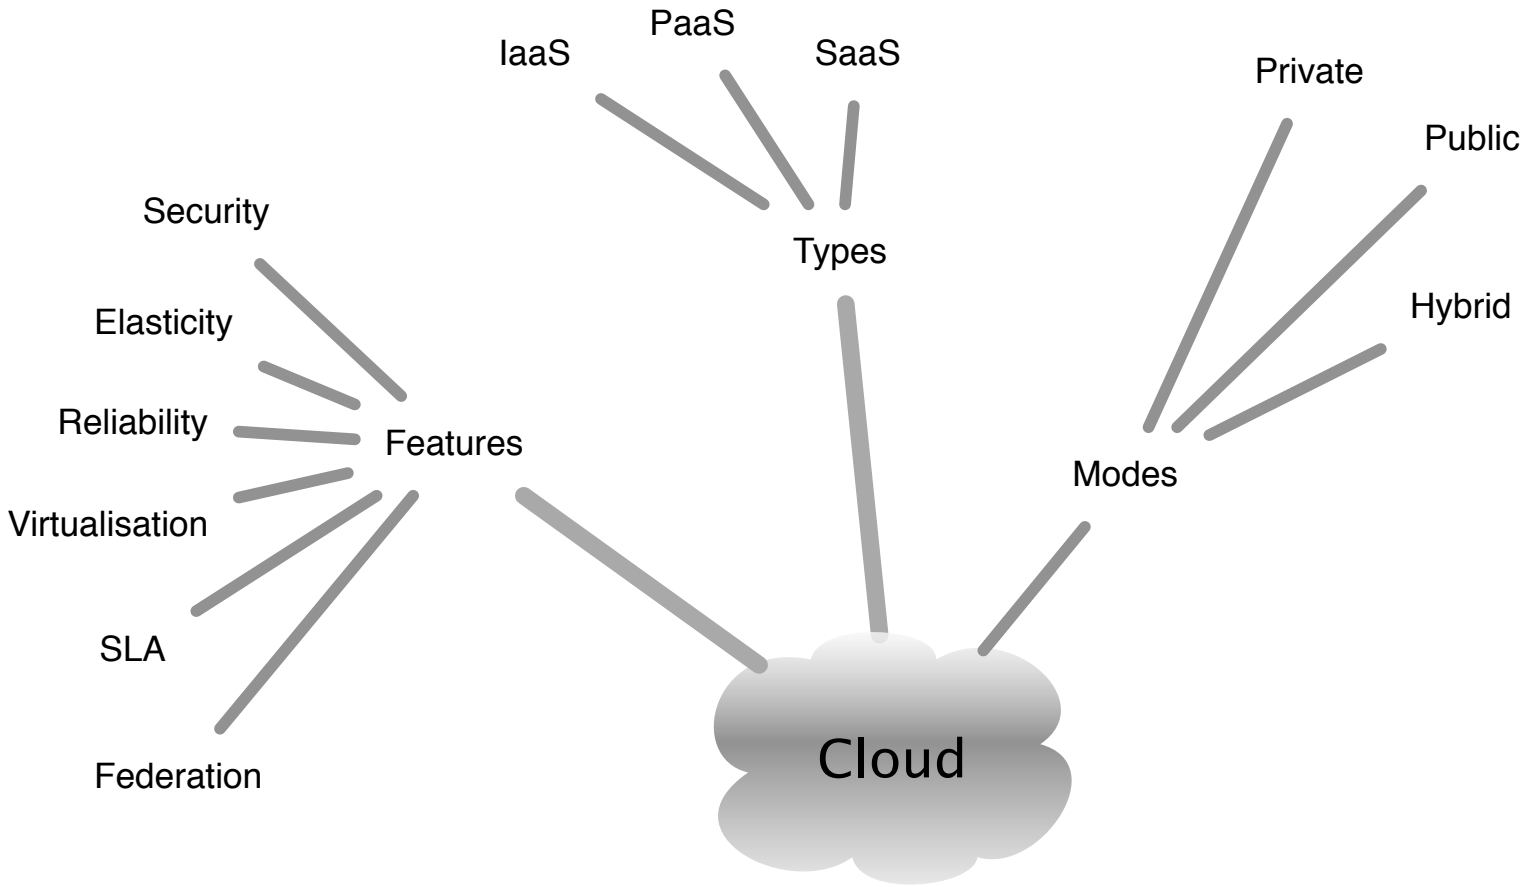

http://cordis.europa.eu/fp7/ict/ssai/docs/cloud-report-final.pdf

Infrastructure as a Service

- Get access on demand to a large number of highly virtualized resources
  - Dynamicity, elasticity
- Concept of OS Virtualization
  - OS does not matter anymore !
  - OS are just software libraries and does not play a central role!
  - Concept of virtual machines to host instances of OS
  - Physical resources are shared by several virtual machines

Properties:
- Isolation
- VM portability
- Suspend/restart

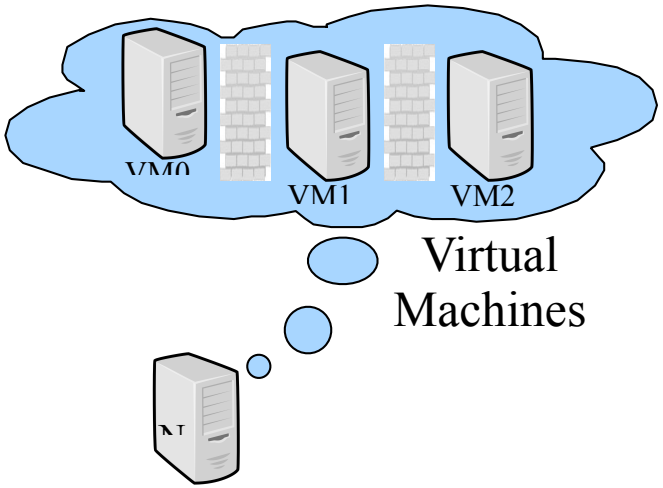

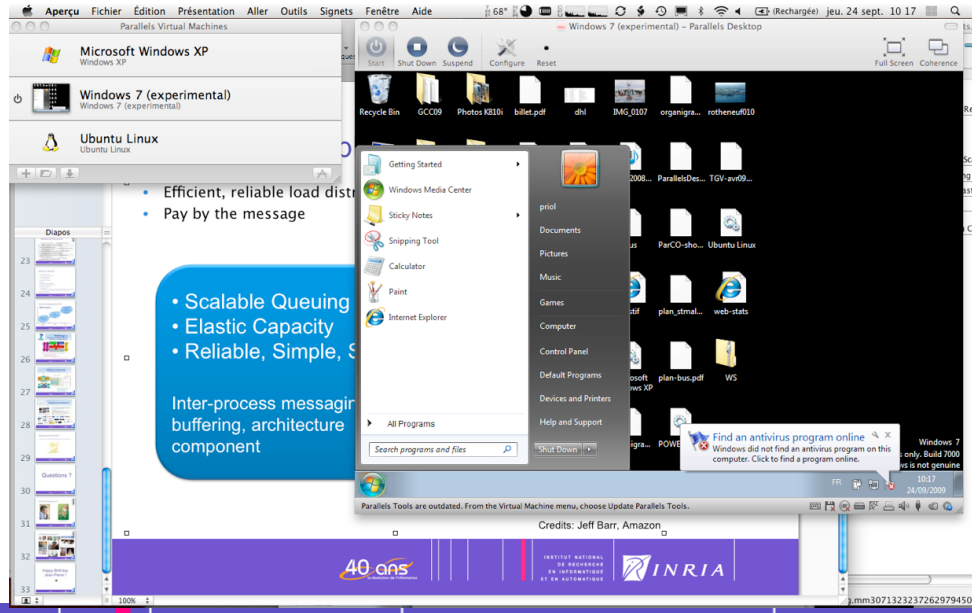

Let's take an example... Amazon !

COMPUTE

Provides on-demand processing
Virtual machine images
pay per server hour

Elastic Compute Cloud Service

MESSAGING

Efficient, reliable comm. layer
Pay by the message

Simple Queue Service

STORAGE

Virtually infinite storage capacity
Objects from 1 byte to 5 gigabytes of data each
pay per GB-month

Simple Storage Service

to create storage volumes
from 1 GB to 1 TB
pay per GB-month

Database service

highly available, scalable, and flexible
non-relational data store
pay per hour

Amazon Pricing - 2010

US – N. Virginia | US – N. California | EU – Ireland

Standard On-Demand Instances	Linux/UNIX Usage	Windows Usage
Small (Default)	$0.085 per hour	$0.12 per hour
Large	$0.34 per hour	$0.48 per hour
Extra Large	$0.68 per hour	$0.96 per hour
High-Memory On-Demand Instances	**Linux/UNIX Usage**	**Windows Usage**
Extra Large	$0.50 per hour	$0.62 per hour
Double Extra Large	$1.20 per hour	$1.44 per hour
Quadruple Extra Large	$2.40 per hour	$2.88 per hour
High-CPU On-Demand Instances	**Linux/UNIX Usage**	**Windows Usage**
Medium	$0.17 per hour	$0.29 per hour
Extra Large	$0.68 per hour	$1.16 per hour

* Data Transfer In will be $0.10 per GB after June 30, 2010.

There is no Data Transfer charge between Amazon EC2 and other Amazon Web Services within the same region (i.e. between Amazon EC2 US West and Amazon S3 in US West). Data transferred between Amazon EC2 instances located in different Availability Zones in the same Region will be charged Regional Data Transfer. Data transferred between AWS services in different regions will be charged as Internet Data Transfer on both sides of the transfer.

Data Transfer In	
All Data Transfer	Free through June 30, 2010*

Data Transfer Out	
First 10 TB per Month	$0.15 per GB
Next 40 TB per Month	$0.11 per GB
Next 100TB per Month	$0.09 per GB
Over 150 TB per Month	$0.08 per GB

Amazon Pricing - 2010

US – N. Virginia | US – N. California | EU – Ireland

Amazon EBS Volumes

- $0.10 per GB-month of provisioned storage
- $0.10 per 1 million I/O requests

Amazon EBS Snapshots to Amazon S3 (priced the same as Amazon S3)

- $0.15 per GB-month of data stored
- $0.01 per 1,000 PUT requests (when saving a snapshot)
- $0.01 per 10,000 GET requests (when loading a snapshot)

US – N. Virginia | US – N. California | EU – Ireland

Amazon EC2 Monitoring

- $0.015 per instance-hour (or partial hour)

US – N. Virginia | US – N. California | EU – Ireland

- $0.025 per Elastic Load Balancer-hour (or partial hour)
- $0.008 per GB of data processed by an Elastic Load Balancer

Platform as a Service

- An application development, deployment and management fabric.
- User programs web service front end and computational & Data Services
- Framework manages deployment and scale out
- No need to manage VM images

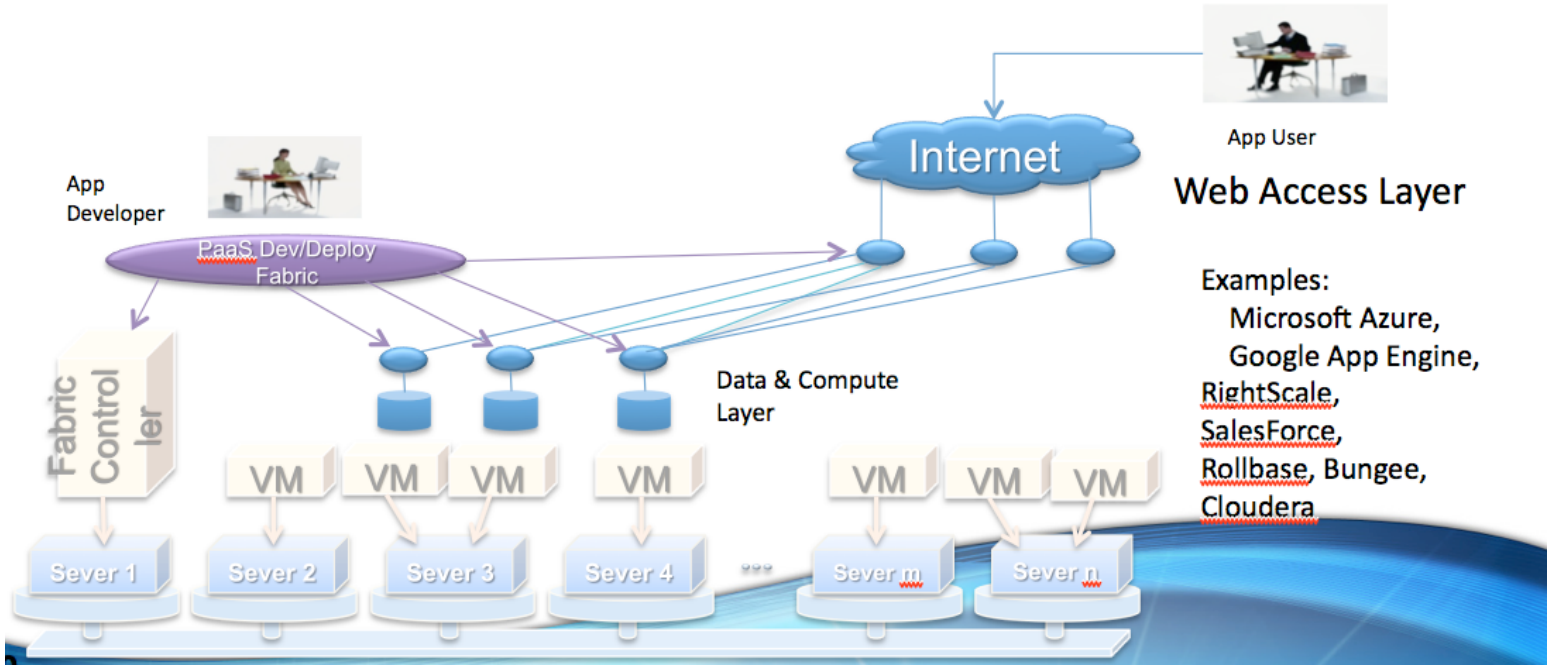

(c)

Software as a Service

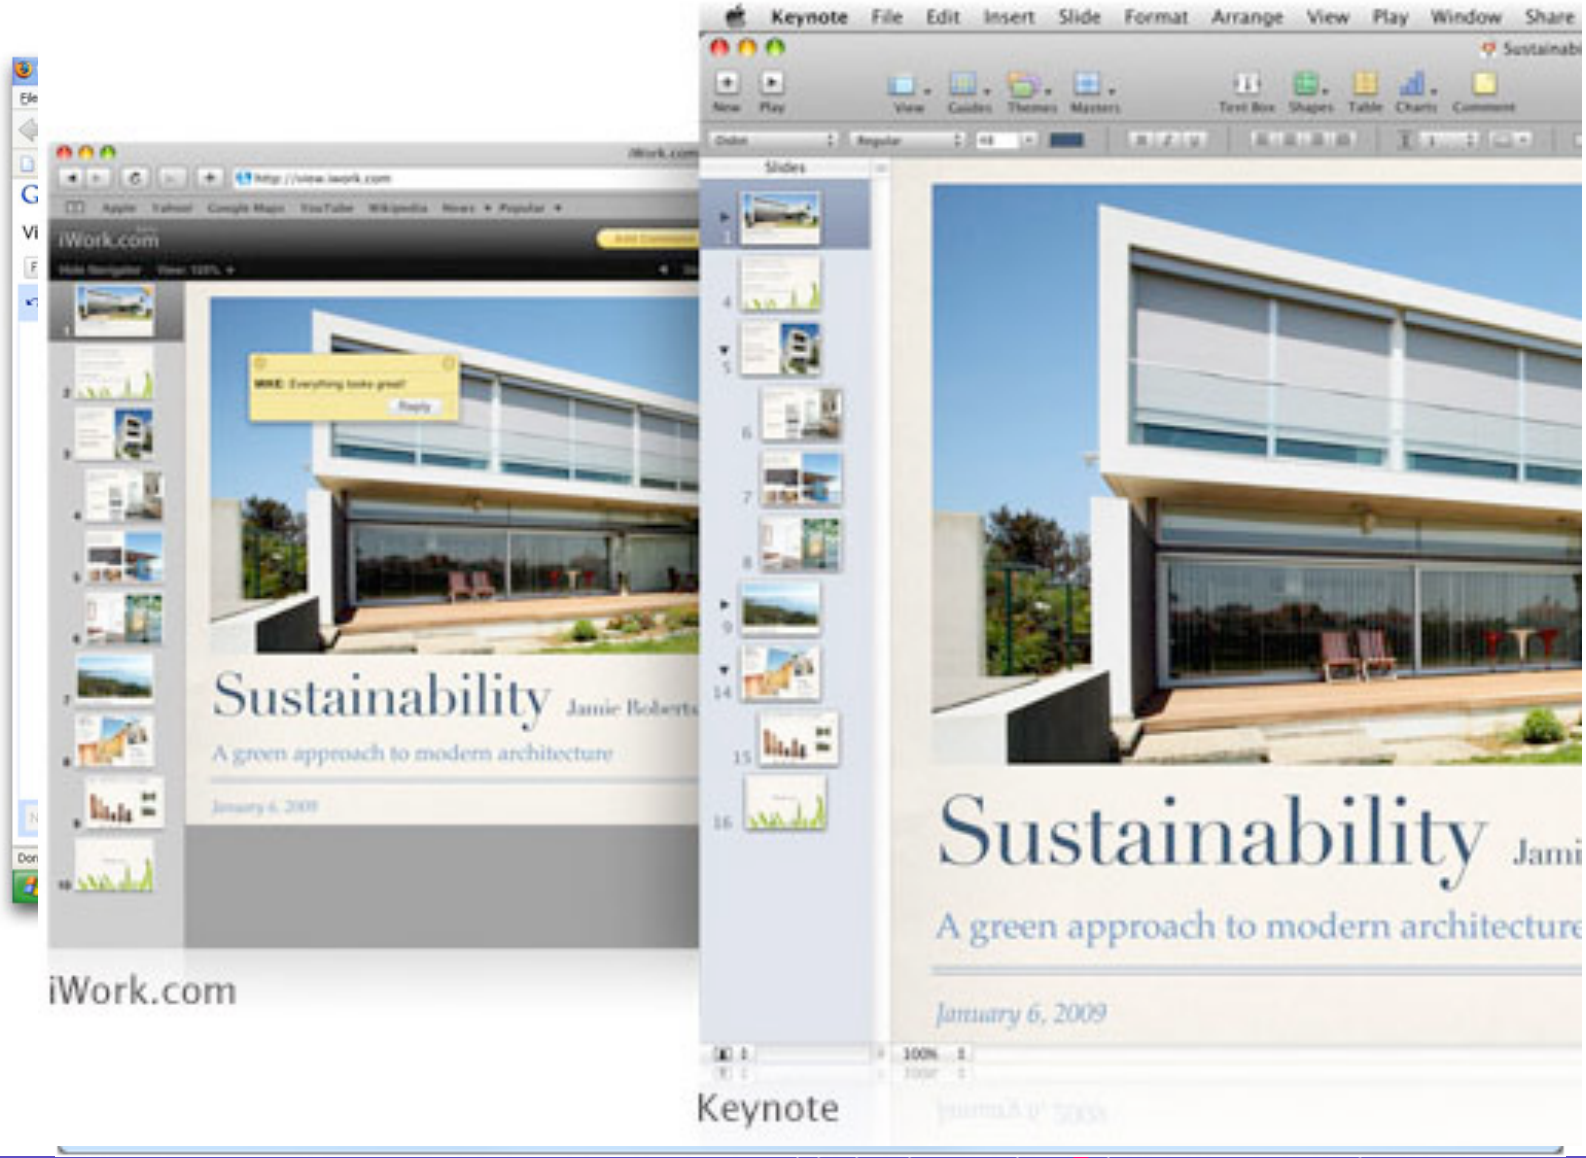

What are the benefits of a SaaS approach

- Avoid managing/installing/deploying new software / patches / update

- Facilitating collaboration between users
  - No more versions to be merged with potential incoherencies

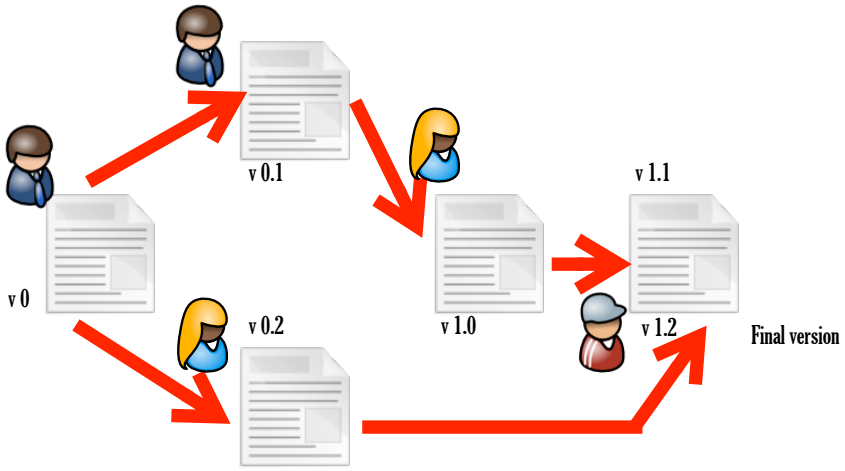

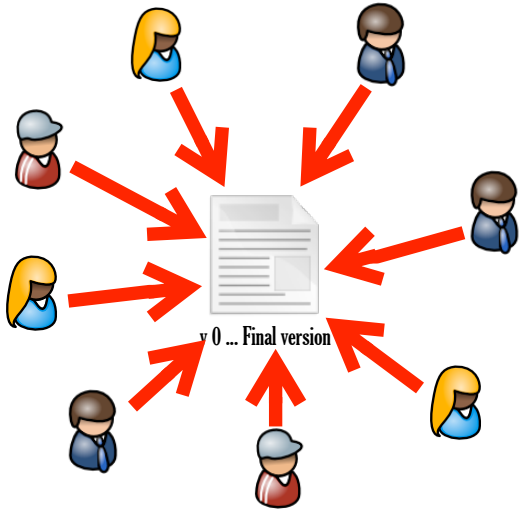

We have only seen the virtuous side !

What is the dark side of Cloud Computing ?

Some research issues with Cloud Computing

- Reliability / Resilience / Fault-tolerance
- Trust, Security and Privacy
- New economical models for computing
- Service Level Agreement / Quality of Service - *From Best Effort to SLA*
- Building cloud-aware applications from legacy applications
- Energy management
- Data management
- Cloud federation
- Autonomic behaviors / Self-*
- Brokering / Scheduling
- Programming models (MapReduce, ...)
- Interactions between legal aspects (laws) and computer science
  - privacy and liability

Reliability / Resilience / Fault-tolerance

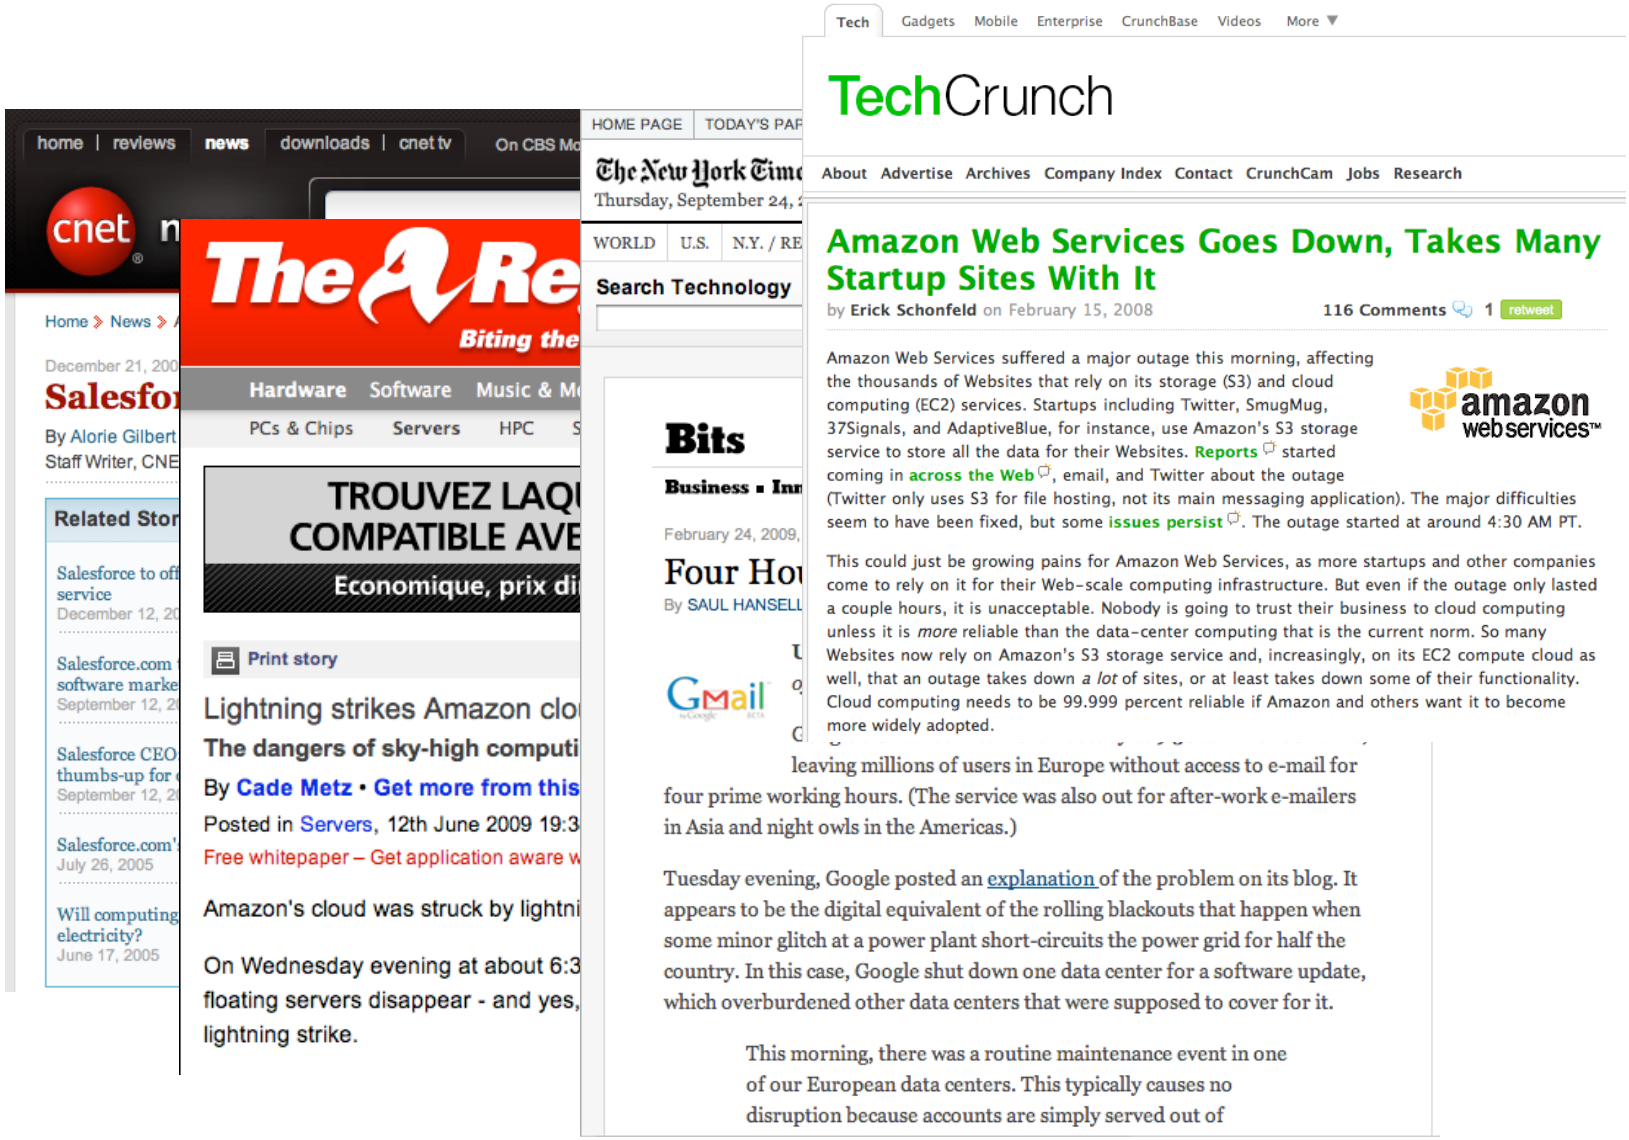

What about failures in the Cloud

- http://www.lemondeinformatique.fr/actualites/lire-les-pannes-dans-le-cloud-ont-coute-71-7-millions-de-dollars-depuis-2007-49375.html

Trust, Security and Privacy

The Apple, Mac, iPod, and iPhone Experts

Macworld

Mac | iPod | iPhone | More Products | News | Reviews

NEW iPhone App Reviews • Mac OS X • Entertainment & HDTV • Create • Business Center • Mobile Mac • Digital

Macworld Ultimate iPhone Guide — Magazine Subscribe & Get a Bonus CD — Customer Service

Macworld » Security

Sep 4, 2009 12:34 am | 0 Comments | + 4

Researchers find a cloud

by Robert McMillan, IDG News Service

Amazon and Microsoft have been pushing outsource raw computing power, but the p have yet to be fully explored, according to and the Massachusetts Institute of Techn

LesEchos.fr

Communication

Jeudi 24 septembre 2009

Météo

Bourse | Finances perso | Les Echos PME | Actualité | Idées & Débats

Actualité | Analyses | Dossiers | Diaporamas | Interviews | En direct des médias du monde | Carnet | Do

Info & Secteurs ▸ Communication ▸ Actualité

HSBC

Partager :

Twitter piraté, Google mis en cause

[17/07/09 - 09H36 - actualisé à 09:41:00]

Le site de micro-blogs, qui utilise les services de dématérialisation info retrouvé toute sa stratégie divulguée sur internet.

Le géant de l'internet Google, promoteur le plus en pointe de l'informatique dé cause depuis que le site de micro-blogs Twitter, client de ses applications pour toute sa stratégie divulguée sur internet.

L'affaire, qui met en émoi la Silicon Valley, a été révélée par le co-fondateur d message au titre humoristique sur le blog de sa société: "Twitter, encore plus souhaitions". Un blogueur français, Korben, et le site spécialisé TechCrunch ont e de documents, et toute la stratégie de Twitter, telle qu'elle se présentait en retrouvée divulguée.

Biz Stone a eu beau dédouaner Google - "Cette attaque n'avait rien à voir av applications Google, que nous continuons d'utiliser" - l'affaire a suscité des dou plus gênantes pour Google que celui-ci multiplie les initiatives pour pénétrer le L'informatique dématérialisée ("cloud computing"), qui consiste pour une entrepri de calcul et de stockage disponible sur l'internet pour réduire ses dépenses en est le nouvel enjeu de la guerre que se livrent Google et les groupe de hautes tec

Autos | Real Estate | Jobs | Classifieds | Mobile | Contests | Place an Ad | DailyNews

NYDailyNews.com

DAILY NEWS

SITE | BLOGS | DISCUSSION

News | Sports | Gossip | Entertainment | NY Events | Local | Opinion | Lifest

News › World › Security experts: China hacked Google to steal U.S.

Article | Comments (2) | Email | Print | RSS | Share

Security experts: China hacked Google to steal U.S. industrial secrets

BY DIANE MOY SCHAEFER
DAILY NEWS STAFF WRITER

Thursday, January 14th 2010, 2:33 PM

Jin/Getty

A Chinese flag flies over the company logo outside the Google China headquarters in Beijing Thursday.

Security experts say Google, along with dozens of other major companies, was the victim of a concerted espionage effort that they say came from within China and exploited flaws in e-mail attachments to get into networks of major financial and technology firms.

At least 34 companies, including Adobe, Symantec, Yahoo and Dow Chemical, were attacked, according to industry sources. On Tuesday, Google revealed that hackers broke into the G-mail accounts of Chinese human rights advocates in the United States, Europe and China, and the search-engine company threatened to pull out of operations in China in response.

RELATED NEWS

ARTICLES

Google may ax China service, end search censoring after hackers launch human rights cyber attack

Google gets even bigger! Investment arm grows partner ranks

Trust, Security and Privacy

- Cloud will introduce new vulnerabilities and threats by allowing a physical infrastructure to be shared thanks to virtualisation technologies
  - The provider is not the only one that could have a malicious behavior...
  - Several VMs from different customers will share the same processor
  - Are we confident that virtualisation can provide 100% isolation across VMs ?
- Have a look at this very interesting paper:
  - ***Hey, You, Get Off of My Cloud: Exploring Information Leakage in Third-Party Compute Clouds***
    Thomas Ristenpart, Eran Tromer†, Hovav Shacham∗, Stefan Savage∗, ∗University of California, †Massachusetts Institute of Technology.
    Published in the proceedings of CCS'09.
  - The paper is about how a cloud customer can «attack» another customer of the same cloud infrastructure
  - It just costs a few $$$ to have a reasonable chance to observe what a cloud user is doing...
  - It has not been fully experimented but the paper gives some indications especially for Amazon EC2
- The threat model
  - Determine where is the VM that hosts a service to be attacked
  - Determine if the attacker's VM co-resides with the VM to be attacked
  - If not, try to launch new VMs until you are co-resident with the VM to be attacked
  - Exploit cross-VM information leakage once co-resident (CPU caches, branch target buffers, network queues, ...)

Virtual Machine instances

- IaaS-based Cloud allows the uploading of virtual machine instances
  - Software for IaaS Clouds tends to be distributed thanks to virtual machine instances (Cloud App Store)
  - Virtual machine instances are prepared/packaged by unaware users

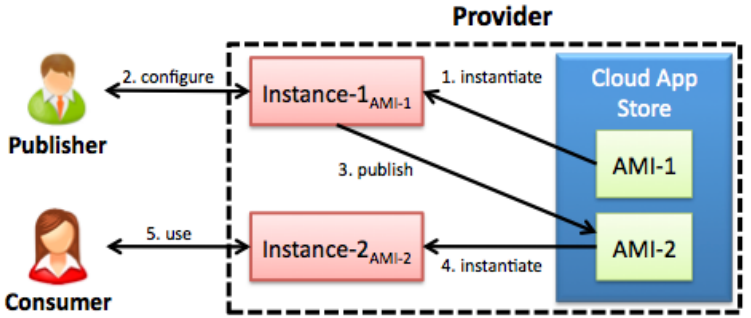

- Have a look at this very interesting paper:
  - ***AmazonIA: When Elasticity Snaps Back***
    Sven Bugiel, Stefan Nürnberger*, Thomas Pöppelmann†, Ahmad-Reza Sadeghi*†, Thomas Schneider*, *TU Darmstadt, †FhG*
    Published in the proceedings of the18th ACM Conference on Computer and Communications Security (CCS'11).
    http://trust.cased.de/AMID
  - The paper is about vulnerabilities associated with the public availability of Amazon Machines Images (AMI) and their deployment using Amazon EC2
  - Highly sensitive information (passwords, keys and credentials) can be extracted from publicly available AMIs
  - 1225 AMIs have been tested letting the authors to get access to source code repositories, administrator passwords, credentials of various web service providers.

Are Cloud infrastructures less secure than non-Cloud ones ?

- **«one of the fastest and easiest ways to access corporate data is through unprotected PDAs that are lost or stolen, as they contain business names and addresses, spreadsheets and other corporate documents»** http://www.theregister.co.uk/2004/09/01/pda_sec
- **«60% of corporate data resides unprotected on PC desktops and laptops»** (IDC analyst Cynthia Doyle, Business Continuity in 2002: It's Not Business as Usual, April 2002)
- Read from http://www.nationalpost.com/
  - **10% of laptop computers will be stolen within the first 12 months of purchase**.
  - **90% of stolen laptops are never recovered.**
  - **49% of companies have had laptops stolen with the last 12 months.**
  - **57% of corporate crimes are linked to stolen laptops.**
  - **80% of computer crime consists of "inside jobs" by disgruntled employees.**
  - **73% of companies had no specific security policies for their laptops in 2003.**
- **66 % of USB thumb drive owners report losing them, over 60 % with private corporate data on them!**

New economic/business model for computing

- Considering a Cloud cost model (such as the Amazon one), what are the impacts on how we design / produce software ?
- Have a look at this very interesting paper:
  - ***The cost of Doing Science on the Cloud: The Montage Example***
    Ewa Deelman, Gurmeet Singh, Miron Livny, Bruce Berriman, John Good, Published in the proceedings of SC'08.
  - The paper is about to find the right balance between cost and performance considering a cost model
  - Based on an astronomy (data-intensive) application (workflow) to deliver on-demand a science-grade mosaic of the sky

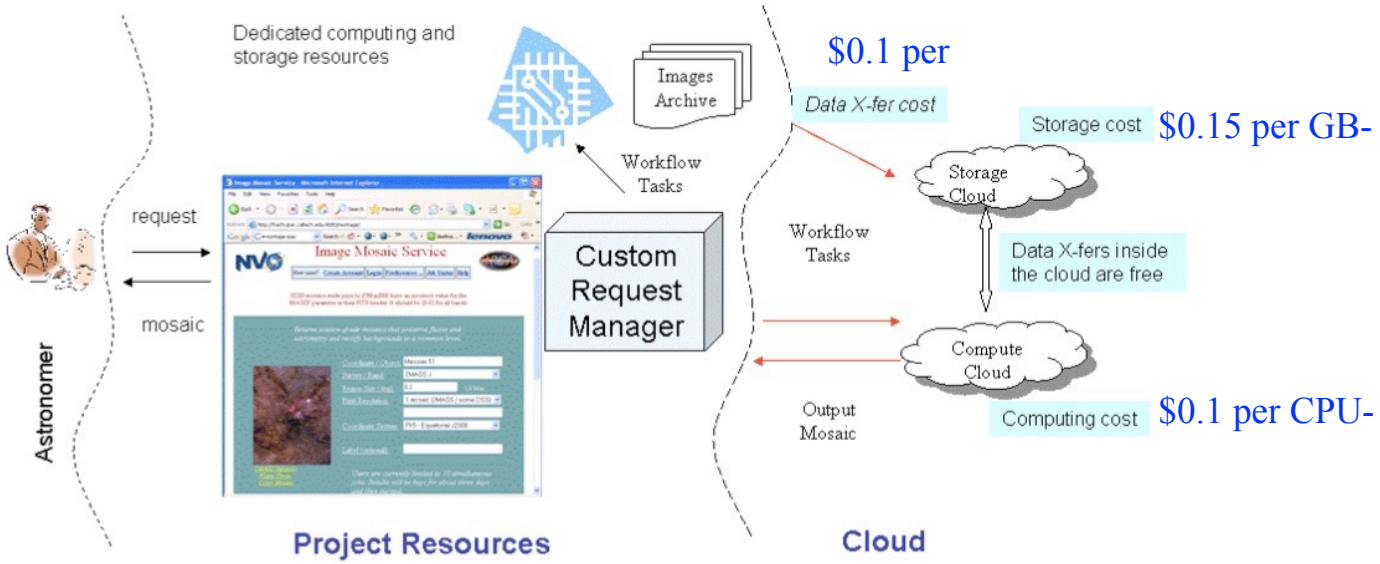

What are the findings ?

- Several implementation data management models
  - Remote I/O : stage in/stage out files at each step of the w
  - Regular: intermediate files produced by the execution of t storage service (S3 for Amazon). Files are deleted when
  - Dynamic cleanup: files are deleted when they have outliv
- How many processors should be used, what will be

	Small	Medium	Large
1 proc	5.5h / 0.60$	20.5h / 2.25$	85h / 9$
128 proc	18mn / 4 $	40mn / 8$	1h / 14$

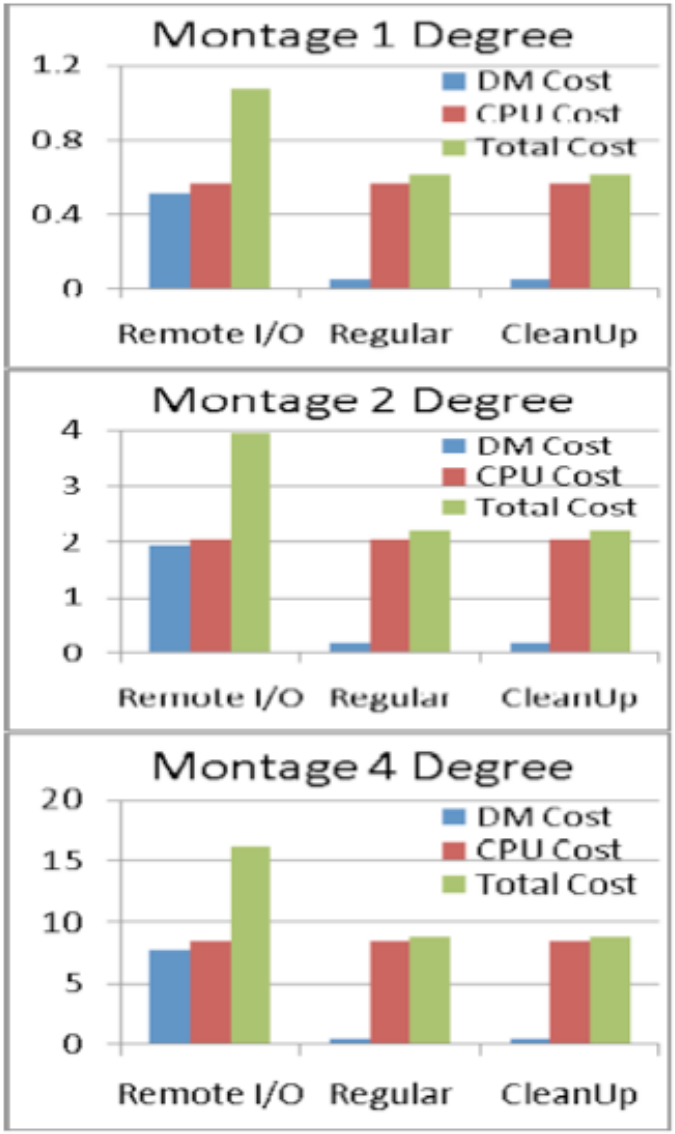

- Does it make sense to archive the generated popul of always generating them on demand from the bas
  - For a small mosaic (173.46 Gbytes), CPU cost to genera
    - For this cost, you can archive it for 21.52 months
  - For a large mosaic (2.229 Tbytes), CPU cost to generate it is $8.40
    - For this cost, you can archive it for 25.12 months
  - Conclusion: if there will be a similar request coming within two years, then it would be cost effective to save popular mosaics of the sky in the cloud...

Conclusions

- Cloud is becoming a buzzword... a lot of hype around it
  - Not the swiss knife for distributed computing (as the grid was supposed to be...)
  - More an evolution than a revolution
  - Less ambitious than Grid but there is an increasing public and business demand
- But there are new opportunities for research:
  - Reliability / Resilience / Fault-tolerance
  - Trust, Security and Privacy
  - New economical models for computing
  - Service Level Agreement / Quality of Service - *From Best Effort to SLA*
  - Building cloud-aware applications from legacy applications
  - Energy management
  - Cloud federation
  - Autonomic behaviors / Self-*
  - Brokering / Scheduling (performance, energy, ...)
  - Programming models (MapReduce, ...)
  - Interactions between legal aspects (laws) and computer science - *privacy and liability*

Questions ?