

# Cloud Computing

Belmiro Moreira

CERN IT-PES-PS

Lecture 2 – May 31, 2012

# Disclaimer



"The Lighter Side Of The Cloud" Comic by CloudTweaks – David Fletcher

# Why Cloud Computing?

- In the 20<sup>th</sup> century:
  - Industrial efficiency and economies of scale based in reliable predictions of customer demand and stable prices for raw materials
  - Large investments in infrastructure to achieve economies of scale
  - All knowledge inside the company
  - High fixed costs to support the IT infrastructure
  - Products with long lifecycles

# Why Cloud Computing?

- In the 21<sup>th</sup> century:
  - Economy volatility is increasing
  - Necessity to accommodate continuous changes in business operations
  - Responsiveness trumps efficiency
  - Companies/organizations need IT infrastructures that enable them to operate efficiently and be responsiveness
  - Other players are specialized in services that companies run internally

# Why Cloud Computing?

- For most companies IT operations are not part of their core competences
- But... they can't launch new products or services or redesign internal operations without new application systems to support them
  - Technology doesn't generate revenue for most companies. Business process powered by technology generates revenue and profits

# Why Cloud Computing?

- Traditional IT infrastructure:
  - Critical in major companies businesses
  - Supports HA
  - About 70 to 80 percent of a company IT budget goes to the operation and maintenance of existing systems
  - Fixed cost model
  - Big up-front capital expense
  - Fixed capacity
  - Resources underused

# Why Cloud Computing?

- Typical IT cost structure:
  - Data Center
  - IT staff
  - Support
  - Software Licenses
  - Operations
  - Maintenance / Development



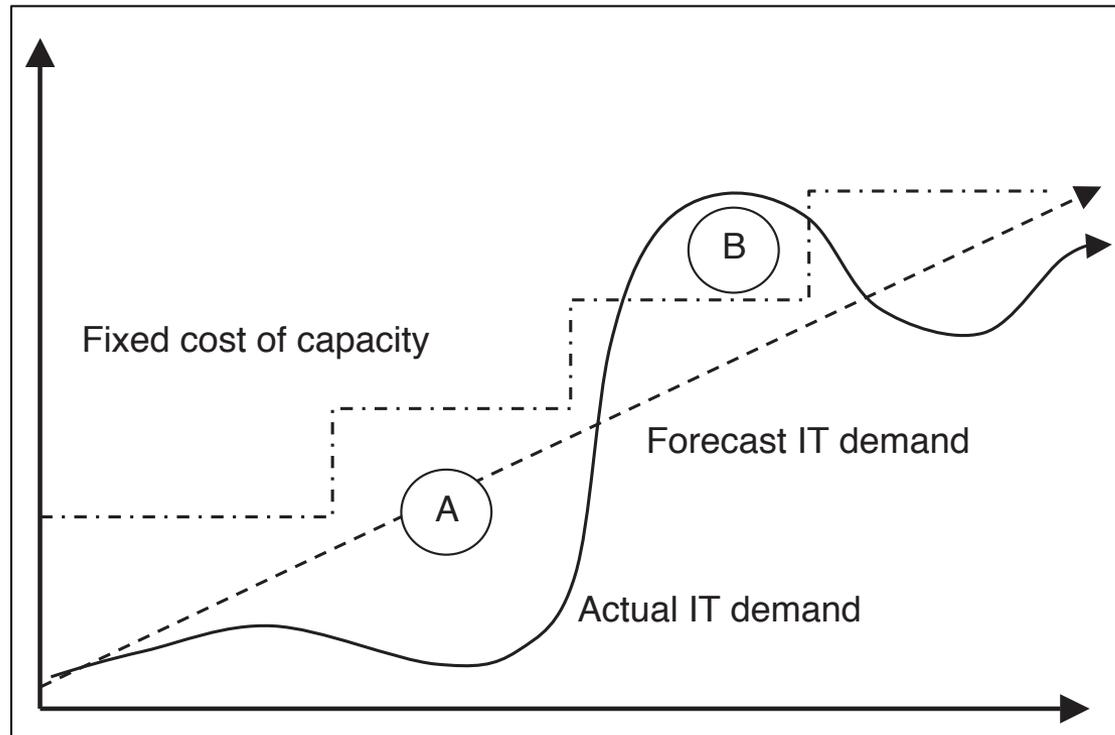


**As business expands... you need more  
computational/storage capacity...**

**Business contraction means... increasingly  
unused computer/storage capacity**

# Why Cloud Computing?

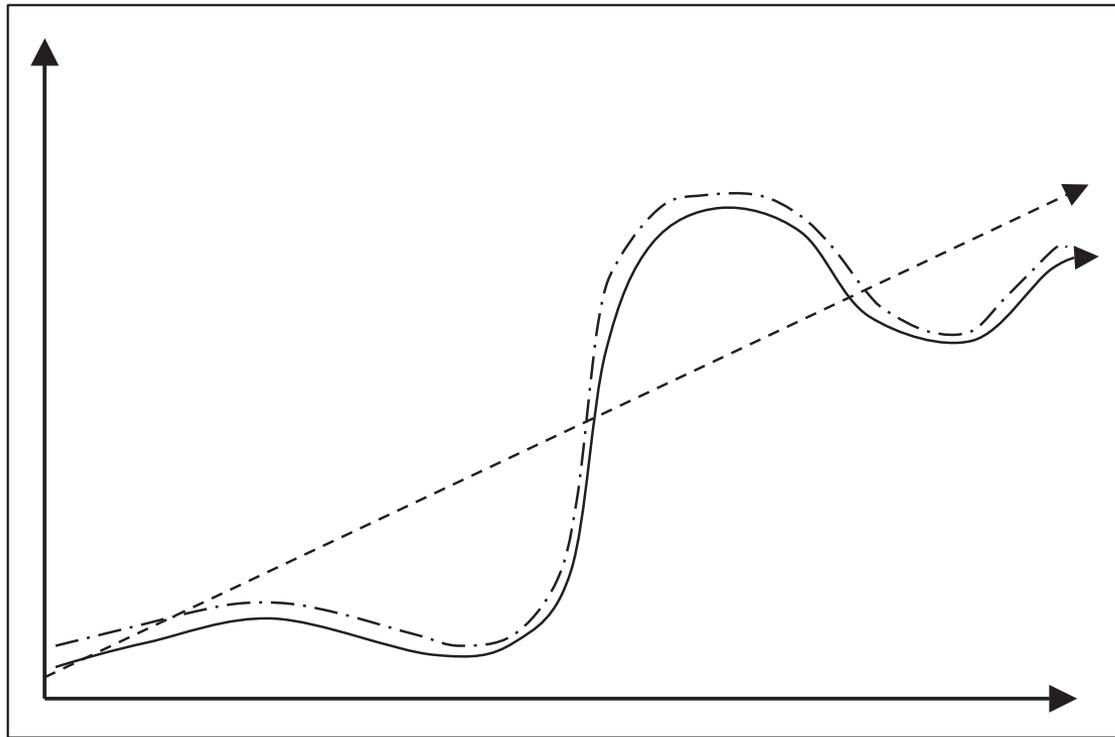
- Fixed cost IT model



"Business in the Cloud" - Michael Hugos, Derek Hultizky

# Why Cloud Computing?

- Variable cost IT model



"Business in the Cloud" - Michael Hugos, Derek Hultizky

# Why Cloud Computing?

- Desirable characteristics for the new IT infrastructure model:
  - Low capital expense
  - Variable cost operations
  - Scalable computing operations
    - Limited computing resources
    - Limited storage

# Why Cloud Computing?

- Previous competitive advantages in business:
  - Electricity
  - Phone / Fax
  - Large Mainframes
  - Desktop computers
  - Internet
  - Virtualization
  - Mobile devices



**Cloud is today's competitive advantage**



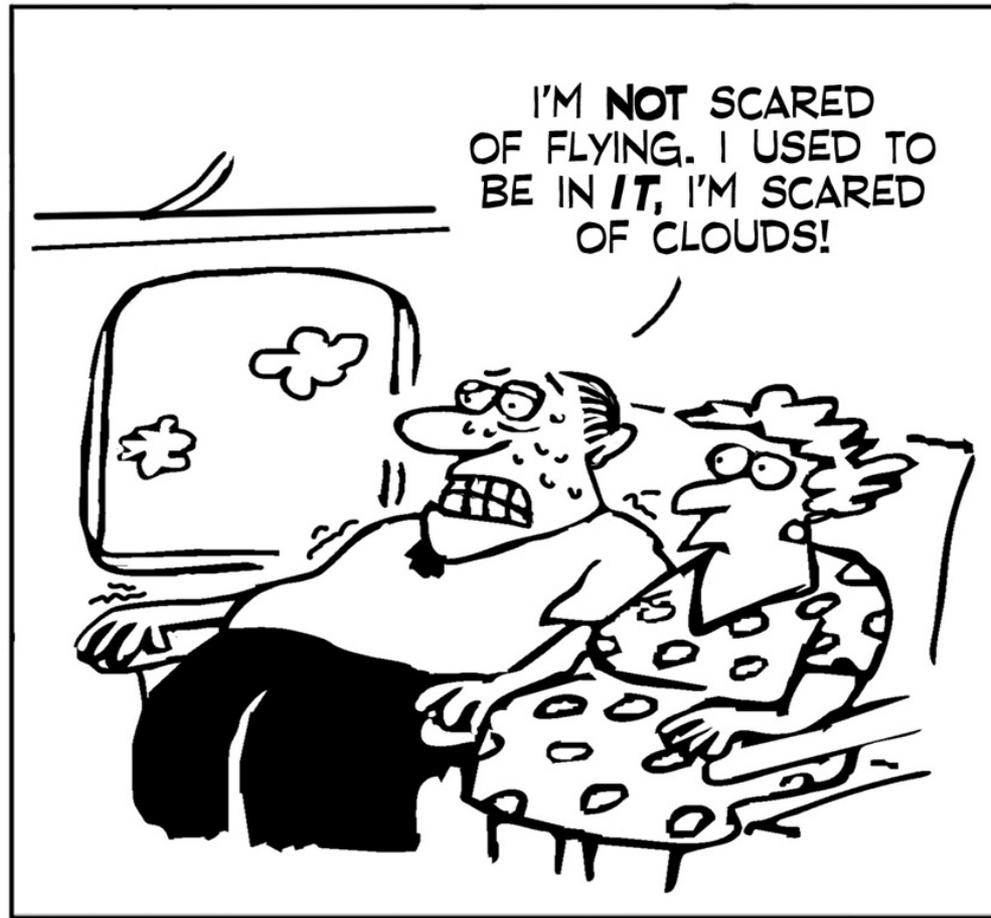
# Why Cloud Computing?

- Main benefits to move into the Cloud
  - Reduced Cost (?)
  - Highly Automated
  - Flexibility
  - Have the capacity when needed
  - Allows IT to shift Focus

# Why Cloud Computing?

- Enables clear focus on the business
- Reduces the capital expense that goes with an internal IT infrastructure
- Automatically scales up and down with business volume
- Allows users to buy only what they consume

# Cloud Computing Risks



"The Lighter Side Of The Cloud" Comic by CloudTweaks – David Fletcher

# Cloud Computing Risks

“Any business leader worried about the security and reliability of their data in the cloud should remember that they’ve been trusting, saving, and storing their personal financial assets in an external, virtual banking”

# Cloud Computing Risks

Suppose you bought a pile of gold bars...



# Cloud Computing Risks

Where should you store them?



Your house



Bank

# Cloud Computing Risks

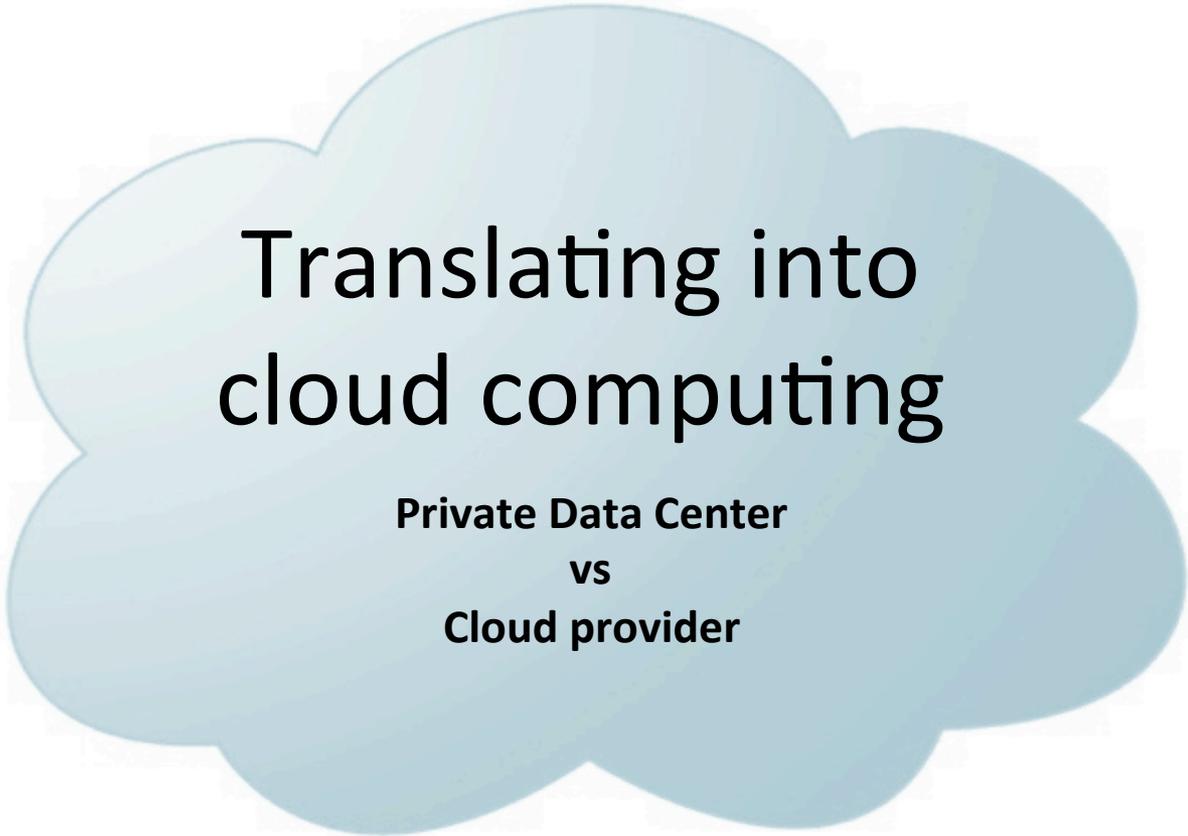
- Thief perspective:
  - Bank
    - Plenty of valuable assets
    - May have an elaborate security protection in place
  - House
    - Some valuable assets
    - May not have an elaborate security protection in place

# Cloud Computing Risks

- House:
  - You can increase the security of your house
    - Locks
    - Fences
    - Alarm
    - Surveillance cameras
    - Access control
      - Family? Friends?
  - What is the cost of private security?

# Cloud Computing Risks

- Bank:
  - Specialized in the protection of valuable assets
  - Security staff
  - Several clients
  - Economy of scale
  - Reputation to preserve



# Translating into cloud computing

**Private Data Center  
vs  
Cloud provider**

# Cloud Computing Risks

- Using public Clouds your compute and storage resources are outside the company
- “Is Cloud Computing secure?”
  - the use of the internet makes it open to attackers worldwide
    - are isolated data centers immune to attacks?
  - where is my data?

# Cloud Computing Risks

- “Can I trust my cloud provider?”
  - With so many companies offering “cloud services” care is needed choosing the cloud provider
  - But, is this process different from selecting a vendor or a provider of any service?

# Cloud Computing Risks

- How to choose a Cloud provider?



[www.cloudtweaks.com](http://www.cloudtweaks.com) – David Fletcher

# Cloud Computing Risks

- How to choose a Cloud provider?
  - Will your cloud service provider be here next year?
    - High reputable Cloud providers are very well established
      - Amazon, Rackspace, HP, IBM, Microsoft, Salesforce.com, ...
    - Very large number of small players with lot of potential
  - Services provided
  - Number of clients and their reputation
  - Previous outages
  - SLA

# Cloud Computing Risks

- Data Loss
  - Most of the Cloud providers replicate the data several times
  - Data can be stored in different availability zones / regions
  - SLAs are very aggressive (ex. 99,9%)
- Data Privacy
  - Encryption

# Cloud Computing Risks

- Data Location
  - Do you know in which country is your data?
  - Your data complies with the country regulations where is located?
  - Network transfer rates
  - Can local authorities search your data?

# Cloud Computing Risks

- SLA – Service Level Agreement
  - Agreement between two parties, customer and service provider
  - Used to refer to the contracted delivery time of the service and/or performance
  - Measuring, monitoring and reporting on cloud performance is based upon an end user experience or the end users ability to consume resources

# Cloud Computing Risks

- Amazon EC2 SLA

- <http://aws.amazon.com/ec2-sla/>

(...) “If the Annual Uptime Percentage for a customer drops below 99.95% for the Service Year, that customer is eligible to receive a Service Credit equal to 10% of their bill “ (...)

- Amazon S3 SLA

- <http://aws.amazon.com/s3-sla/>

Monthly Uptime Percentage	Service Credit Percentage
Equal to or greater than 99% but less than 99.9%	10%
less than 99%	25%



# Recent Cloud outages

# Cloud Computing Risks

- April 21, 2011 - Amazon EC2 outage
  - Some headlines:
    - Amazon's Trouble Raises Cloud Computing Doubts  
*The New YorkTimes*
    - Amazon Outage Crashes Reddit, Quora, and Other Websites  
*PCWorld*

# Cloud Computing Risks

- April 21, 2011 - Amazon EC2 outage
  - Amazon released a summary about the incident:
    - <http://aws.amazon.com/message/65648/>
  - a significant fraction of the volumes concluded that the replication mirroring was out-of-sync and started re-replicating causing further havoc, including an overload of the EBS control plane and network

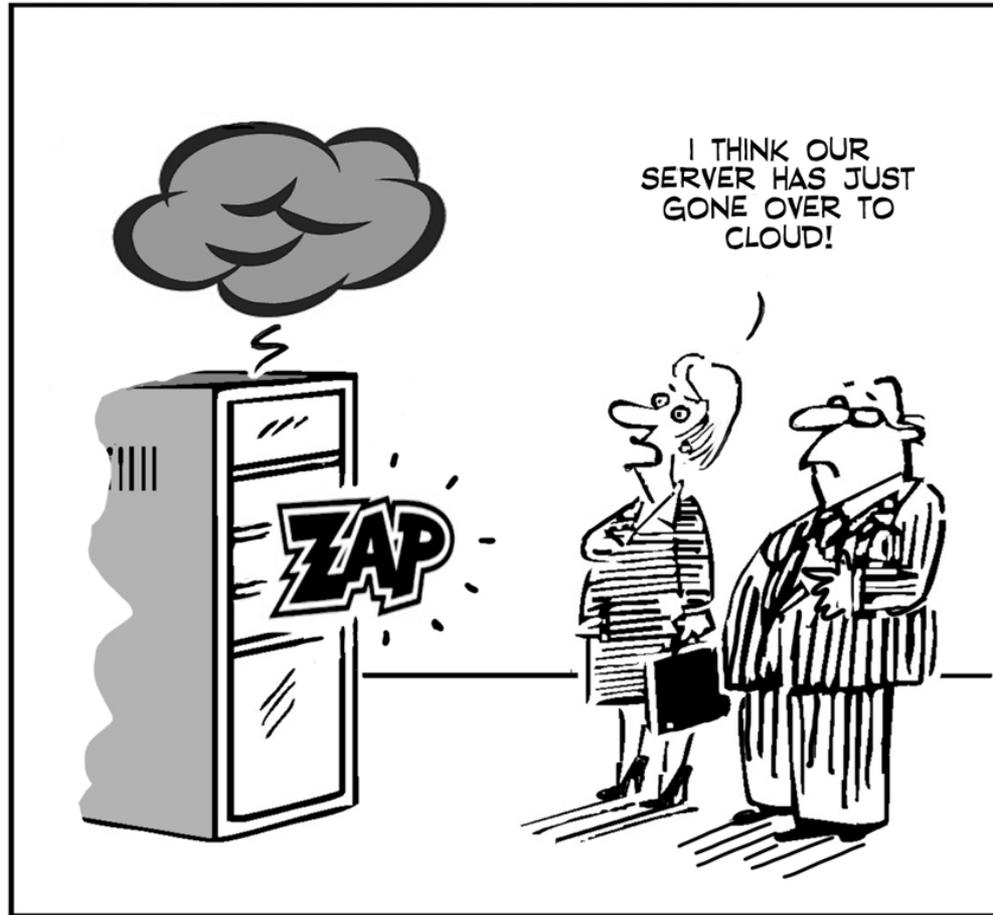
# Cloud Computing Risks

- February 29, 2012 - Microsoft Azure outage
  - A process meant to detect failed hardware in Microsoft's Azure cloud was inadvertently triggered by a Leap Day software bug that set invalid expiration dates for security certificates
  - Service impact of 8-10 hours for users of Azure data centers in Dublin, Ireland, Chicago, and San Antonio
  - The software bug itself was isolated within 2.5 hours and corrected within 7.5 hours of the first incident

# Cloud Computing Risks

- Cloud computing risks:
  - Security
  - Compatibility
  - Availability
  - Monitoring
  - Lock in
  - Standardization

# Migrate to the Cloud



"The Lighter Side Of The Cloud" Comic by CloudTweaks – David Fletcher

# Migrate to the Cloud

- Strategy based on agility and responsiveness
  - Ability to make continuous changes and adjustments in operating procedures, resources, ...
- Continuous exploration of new opportunities and projects
- Reduce initial IT investments
- IaaS, PaaS or SaaS levels

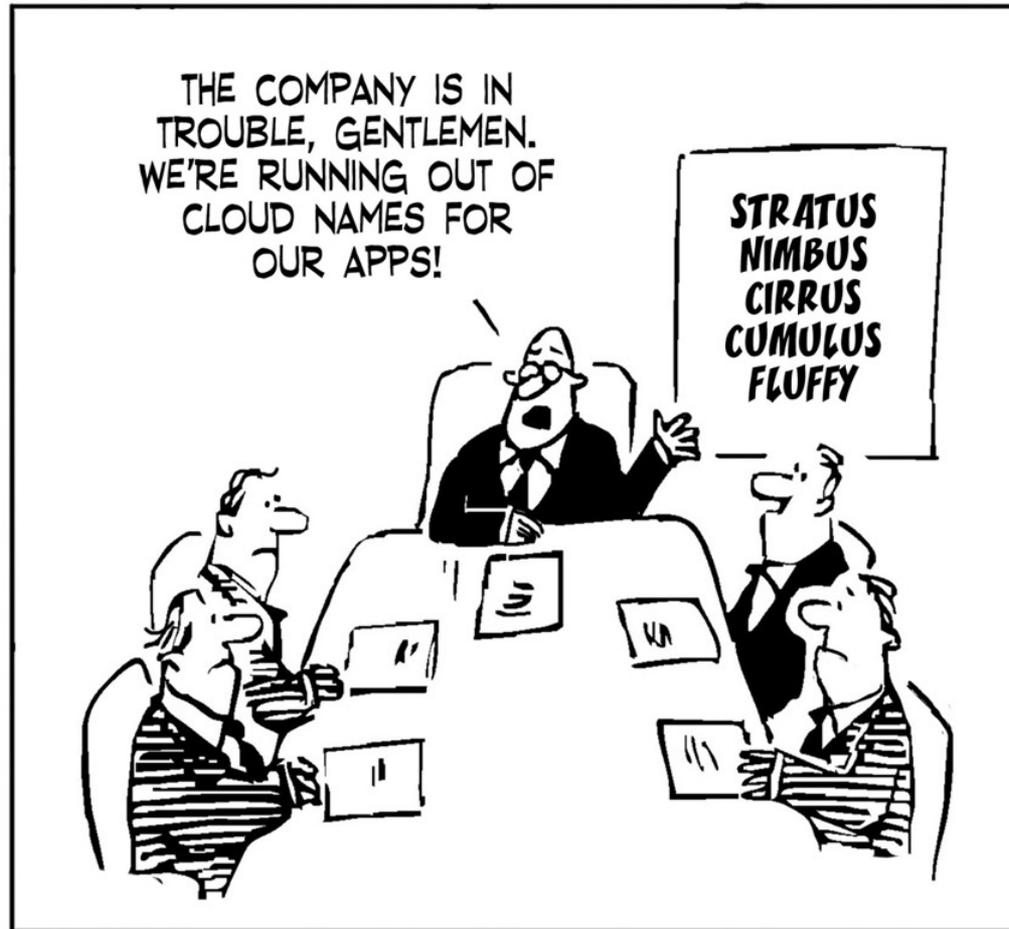
# Migrate to the Cloud

- Private Cloud
  - Companies that already have large investments in data centers
  - Generally an IaaS model
    - Select the right Cloud manager tool
  - Opportunity to implement business specifics not available in public Cloud providers
  - Be agile and improve infrastructure management
  - Reduce costs
  - Applications need to be cloud compatible

# Migrate to the Cloud

- Public Cloud
  - Select one or several Cloud providers
  - Burst into the cloud to have temporary capacity
    - Make sure the applications are designed for the Cloud
  - Eliminate complex license systems if using SaaS
  - Increase development consistency if using PaaS
    - Simplify application integration
    - Common building blocks
  - Start-ups

# Different Cloud Solutions



"The Lighter Side Of The Cloud" Comic by CloudIweaks – David Fletcher

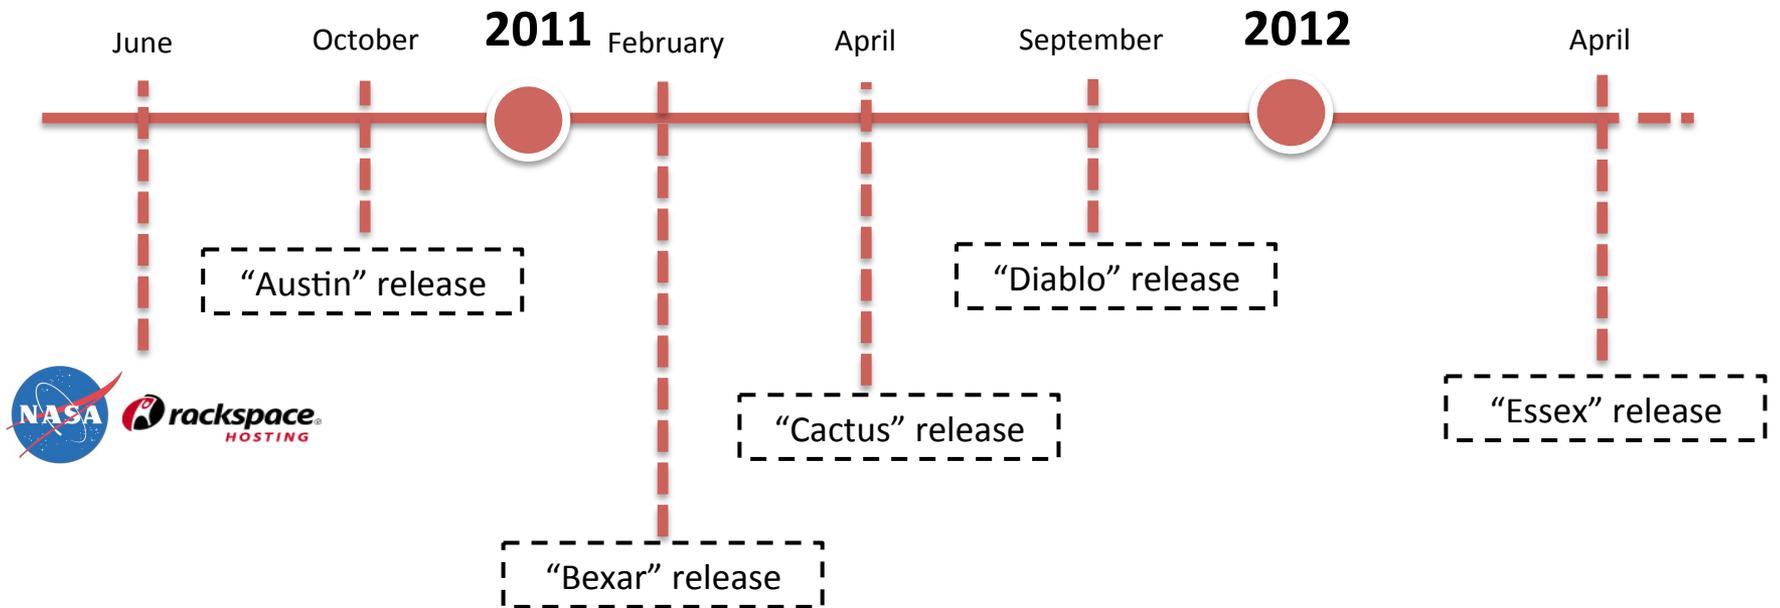
# Openstack

- Openstack is:
  - Community
  - Project
  - Stack of open source software to help organizations run clouds for virtual computing and/or storage.
- Openstack provides an operating platform, or toolkit, for orchestrating clouds.

# Openstack

- Open source software
  - Apache 2 license
- Written in Python
  - Easy to add new functionalities
- No “enterprise” version
  - All source code is available
- Open design process

# Openstack



# Openstack

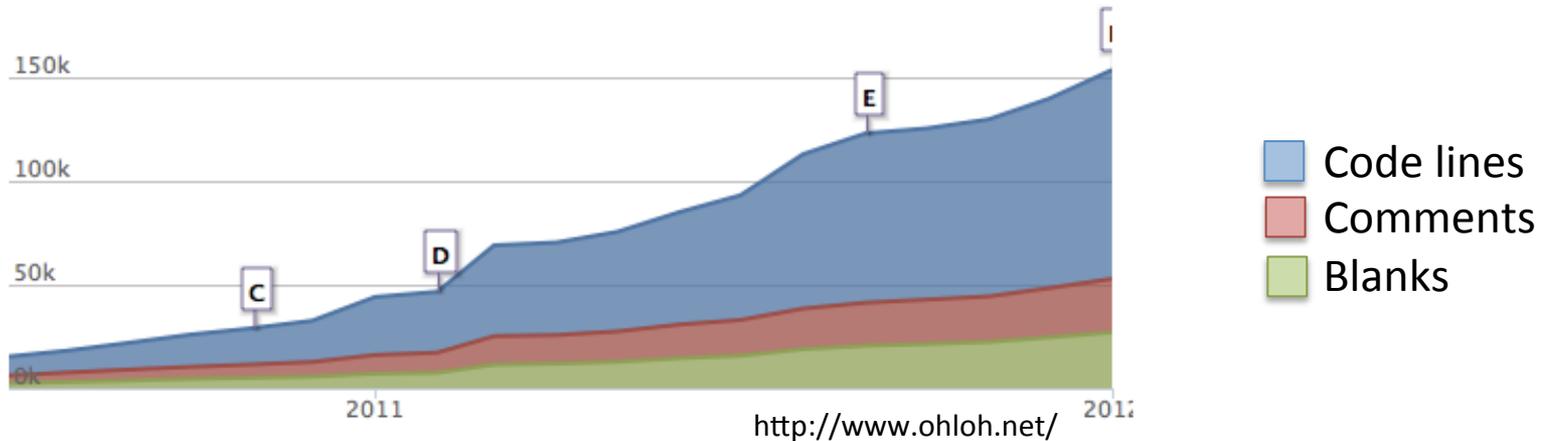
- 176 companies supporting Openstack



- Find the complete list at:
  - <http://www.openstack.org/community/companies/>

# Openstack

- Active community
  - Number of code lines in Openstack Nova



- Regular meetings, conferences and design summits
- Several “user groups” around the world

# Openstack

- Core projects
  - Compute (code named “Nova”)
    - large-scale deployments of automatically provisioned virtual compute instances
  - Storage (code named “Swift”)
    - large-scale, redundant storage of static objects
  - Image service (code named “Glance”)
    - discovery, registration, and delivery services for virtual disk images
  - Identity (code named “Keystone”)
    - unified authentication across all OpenStack projects
  - Dashboard (code named “Horizon”)
    - access and provision cloud-based resources through a self-service portal

# Openstack

- Community projects

- Crowbar

- a DevOps inspired cloud installation and maintenance system that allows users to quickly deploy a fully functioning OpenStack cloud

- Quantum

- provides network connectivity-as-a-service for devices managed by other OpenStack services. It exposes a generic and extensible API, allowing users to build and manage their networks.

- Melange

- provide network information services with a focus on IP address management and address discovery

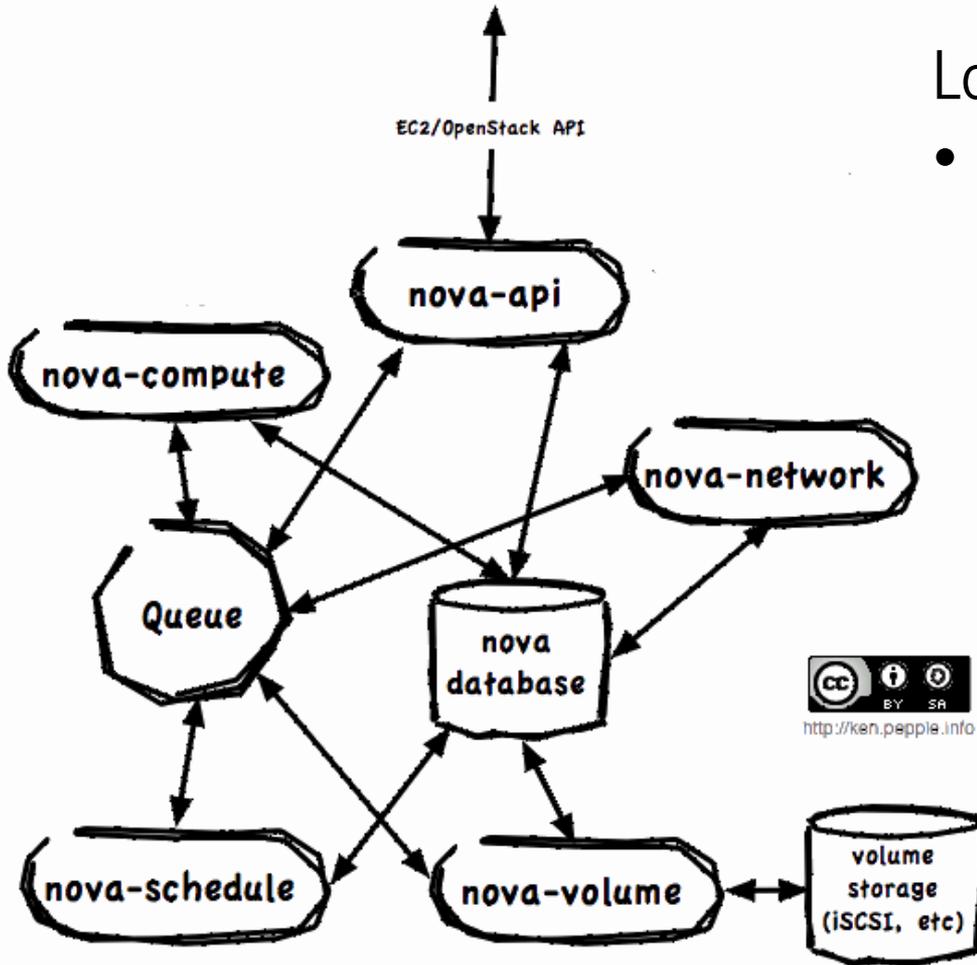
# Openstack

- Openstack Nova
  - Openstack nova is the tool that orchestrate a cloud, including running instances, managing networks, and controlling access to the cloud through users and projects
  - IaaS – Infrastructure as a Service
    - Build a private/public cloud similar to Amazon EC2 and RackSpace Cloud Servers
  - Supports various virtualization technologies:
    - XEN, KVM, UML, QEMU, LXC, VMWare ESX/ESXi

# Openstack

## Logical Architecture (Nova)

- Asynchronous orchestration and information sharing:
  - The queue provides a central hub for passing messages between daemons
    - RabbitMQ
  - The SQL database stores the state of the cloud infrastructure
    - MySQL, PostgreSQL, sqlite



# Openstack

- Next release “Folsom”
  - Goals:
    - minimizing scope
    - block storage breakout
    - network breakout
    - stable trunk
    - metadata/config drive
    - AWS compatibility
    - consistent CLIs
    - operation support
    - security improvements
    - distributed state management
    - cells/host aggregates
    - migrations

# Amazon EC2

- Amazon EC2 – “Elastic Compute Cloud”
  - Announced a limited public beta on August 2006
  - Allows users to rent virtual computers on which to run their own computer applications
  - A user can boot an Amazon Machine Image to create a virtual machine, which Amazon calls an "instance"
  - A user can create, launch, and terminate server instances as needed, paying for active servers on an hourly basis

# Amazon EC2

Region: <input type="text" value="EU (Ireland)"/>		
	Linux/UNIX Usage	Windows Usage
<b>Standard On-Demand Instances</b>		
Small (Default)	\$0.090 per Hour	\$0.115 per Hour
Medium	\$0.180 per Hour	\$0.230 per Hour
Large	\$0.360 per Hour	\$0.460 per Hour
Extra Large	\$0.720 per Hour	\$0.920 per Hour
<b>Micro On-Demand Instances</b>		
Micro	\$0.025 per Hour	\$0.035 per Hour
<b>High-Memory On-Demand Instances</b>		
Extra Large	\$0.506 per Hour	\$0.570 per Hour
Double Extra Large	\$1.012 per Hour	\$1.140 per Hour
Quadruple Extra Large	\$2.024 per Hour	\$2.280 per Hour
<b>High-CPU On-Demand Instances</b>		
Medium	\$0.186 per Hour	\$0.285 per Hour
Extra Large	\$0.744 per Hour	\$1.140 per Hour
<b>Cluster Compute Instances</b>		
Quadruple Extra Large	N/A*	N/A*
<b>Cluster GPU Instances</b>		
Quadruple Extra Large	N/A*	N/A*
* Cluster Compute and Cluster GPU Instances are currently only available in the US East (Virginia) Region.		

Amazon EC2 price table - <http://aws.amazon.com/ec2/pricing/> - May 1, 2012

# Amazon EC2

The screenshot displays the Amazon EC2 Console Dashboard. At the top, the navigation bar includes the AWS Management Console logo, the URL [aws.amazon.com](http://aws.amazon.com), and links for AWS, Products, Developers, Community, Support, and Account. The user is identified as **Welcome, Belmiro** with links for Settings and Sign Out.

The main navigation pane on the left lists various services: Elastic Beanstalk, Amazon S3, Amazon EC2 (selected), Amazon VPC, Amazon CloudWatch, Amazon Elastic MapReduce, Amazon CloudFront, Amazon RDS, and Amazon SNS. Below this, a 'Navigation' section shows the 'Region' set to 'US East' and a list of categories: INSTANCES (with links to EC2 Dashboard, Instances, and Spot Requests), IMAGES (with links to AMIs and Bundle Tasks), ELASTIC BLOCK STORE (with links to Volumes and Snapshots), and NETWORKING & SECURITY (with links to Elastic IPs, Security Groups, Placement Groups, Load Balancers, and Key Pairs).

The main content area is titled 'Amazon EC2 Console Dashboard' and is divided into several sections:

- Getting Started:** A yellow box contains the text: "To start using Amazon EC2 you will want to launch a virtual server, known as an Amazon EC2 instance." Below this is a prominent 'Launch Instance' button. A note states: "Note: Your instances will launch in the US East (Virginia) region."
- Service Health:** A table shows the current status of Amazon EC2 in the US East - N. Virginia region. The status is 'Service is operating normally' with a green checkmark icon. A link is provided to 'View complete service health details'.
- My Resources:** A summary of resources in the US East (Virginia) region, with a 'Refresh' button. The counts are:
  - 0 Running Instances
  - 0 Elastic IPs
  - 0 EBS Volumes
  - 0 EBS Snapshots
  - 1 Key Pair
  - 2 Security Groups
  - 0 Load Balancers
  - 0 Placement Groups
- Related Links:** A list of links including Documentation, All EC2 Resources, Forums, Feedback, and Report an Issue.

# Amazon EC2

- How to create an instance in Amazon EC2

```
ec2-describe-images -o amazon
```

```
IMAGE      ami-225fba4b  ec2-public-images/fedora-core4-apache-mysql-  
v1.07.manifest.xml  amazon      available public          i386 machine  
instance-store paravirtual  
IMAGE      ami-25b6534c  ec2-public-images/fedora-core4-apache-mysql.manifest.xml  
amazon      available public          i386 machine          instance-store  
paravirtual  
IMAGE      ami-2e5fba47  ec2-public-images/fedora-core4-apache-v1.07.manifest.xml  
amazon      available public          i386 machine          instance-store  
paravirtual  
IMAGE      ami-23b6534a  ec2-public-images/fedora-core4-apache.manifest.xml  
amazon      available public          i386 machine          instance-store  
paravirtual
```

```
(...)
```

# Amazon EC2

- How to create an instance in Amazon EC2

```
ec2-run-instances ami-ccf405a5 -k ec2-keypair -t m1.small
```

```
RESERVATION    r-6cb76c01    980927233206    defaultINSTANCE    i-8de9aae1    ami-ccf405a5  
                pending      ec2-keypair     0                m1.small  
2011-02-17T08:09:38+0000 us-east-1c    aki-407d9529    monitoring-disabled  
                ebs                paravirtual
```

```
ec2-describe-instances
```

```
RESERVATION    r-6cb76c01    980927233206    defaultINSTANCE    i-8de9aae1    ami-ccf405a5  
ec2-50-16-58-182.compute-1.amazonaws.com  
ip-10-203-55-143.ec2.internal running      ec2-keypair     0                m1.small  
2011-02-17T08:09:38+0000us-east-1c aki-407d9529    monitoring-disabled  
50.16.58.182    10.203.55.143    ebs                paravirtual
```

# Amazon EC2

- How to create an instance in Amazon EC2

```
ssh -i <where is your keypair> ec2-user@ ec2-50-16-58-182.compute-1.amazonaws.com
```

```
Linux ip-10-203-55-143 2.6.35-24-virtual #42-Ubuntu SMP Thu Dec 2 05:01:52 UTC 2010
i686 GNU/Linux
Ubuntu 10.10
Welcome to Ubuntu!

ubuntu@ip-10-203-55-143:~$
```

```
ec2-terminate-instances i-8de9aae1
INSTANCE i-8de9aae1      running  shutting-down
```

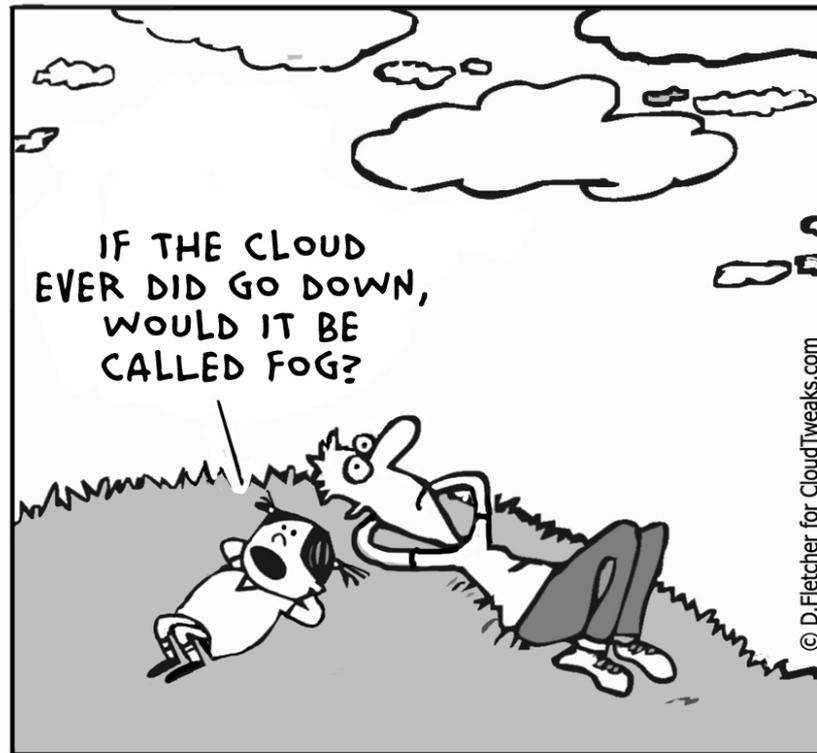
```
ec2-describe-instances
RESERVATION    r-6cb76c01      980927233206  defaultINSTANCE  i-8de9aae1      ami-
ccf405a5      terminated      ec2-keypair   0                m1.small
                2011-02-17T08:09:38+0000 us-east-1c    aki-407d9529    monitoring-disabled
                ebs                paravirtual
```

# References

- “The Big Switch – rewiring the world, from Edison to Google” – Nicholas Carr
- “Business in the Cloud” – Michael Hugos, Derek Hultitzky
- “Cloud Computing – Principles and Paradigms” – edited by Rajkumar Buyya, James Broberg, Andrzej Goscinski
- Sotomayor, B. and Montero, R.S. and Llorente, I.M. and Foster, I.; “Virtual infrastructure management in private and hybrid clouds”; IEEE Internet Computing; 2009.
- Vaquero, L.M. and Rodero-Merino, L. and Caceres, J. and Lindner, M.; “A break in the clouds: towards a cloud definition”; ACM SIGCOMM Computer Communication Review; 2008.
- Renato J. Figueiredo and Peter A. Dinda and José A. B. Fortes; “A Case For Grid Computing On Virtual Machines”; 23rd International Conference on Distributed Computing Systems; 2003.
- Marshall, P. and Keahey, K. and Freeman, T; “Elastic Site: Using Clouds to Elastically Extend Site Resources”; 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing; 2010.
- Ian Foster, Yong Zhao, Ioan Raicu, Shiyong Lu; "Cloud Computing and Grid Computing 360-Degree Compared"; Grid Computing Environments Workshop; 2008.
- Cloud computing definitions and references: <http://jameskaskade.com/?p=594>
- Cloud Tweaks – Cloud Computing Community: <http://www.cloudtweaks.com>
- All illustrations (cartons) from [www.cloudtweaks.com](http://www.cloudtweaks.com) – David Fletcher
- Wikipedia – “Cloud Computing” - [http://en.wikipedia.org/wiki/Cloud\\_computing](http://en.wikipedia.org/wiki/Cloud_computing)
- “I’m cloud confused” by Handy Arjanto - [http://en.wikipedia.org/wiki/Cloud\\_computing](http://en.wikipedia.org/wiki/Cloud_computing)
- “In cloud we trust” by Handy Arjanto - <http://www.slideshare.net/Gupperts/in-cloud-we-trust>

# Thank you

belmiro.moreira@cern.ch



The Lighter Side of The Cloud - Comic by CloudTweaks - David Fletcher