

Ixcloud

Sebastien Goasguen, Ulrich
Schwickerath, Belmiro Moreira,
Romain Wartel and Ewan Roche

A few words on "whole machine"

- Condor can do it

- <https://condor-wiki.cs.wisc.edu/index.cgi/wiki?p=WholeMachineSlot>
- +RequiresWholeMachine = True
- Requirements = IsWholeMachineSlot

- PBS can do it

- ```
o [sebgoa@user001 ~]$ qsub -I -l nodes=1:ppn=8
o [sebgoa@node0347 ~]$ echo $PBS_NODEFILE
/var/spool/torque/aux//1496074.pbs001.palmetto.clemson.edu
[sebgoa@node0347 ~]$ cat /var/spool/torque/aux//1496074.pbs001.
palmetto.clemson.edu
node0347
node0347
node0347
node0347
node0347
node0347
node0347
node0347
node0347
```

# A few words on "whole machine"

And LSF as well.

```
bsub -n N -span[host=1] myjob
```

Current scheduling policy is ***spread them...***

We would need to ***pack them*** ...to reduce the average waiting time to obtain a full node.

# A few words on "community" cluster

Multiple stakeholders buy nodes

Access guaranteed for their share at all times

**Preemptable** queue gives access to non-owners in an opportunistic manner.

Job framework needs to be able to deal with preempted jobs....

# Outline

- Basics of cloud computing
- Key Features
- Internal Cloud
- Hybrid Cloud
- Public Cloud
- Cloud infrastructure at CERN



# What is Cloud Computing ?



Broadcast Yourself™

[Home](#) [Videos](#) [Channels](#) [Shows](#)

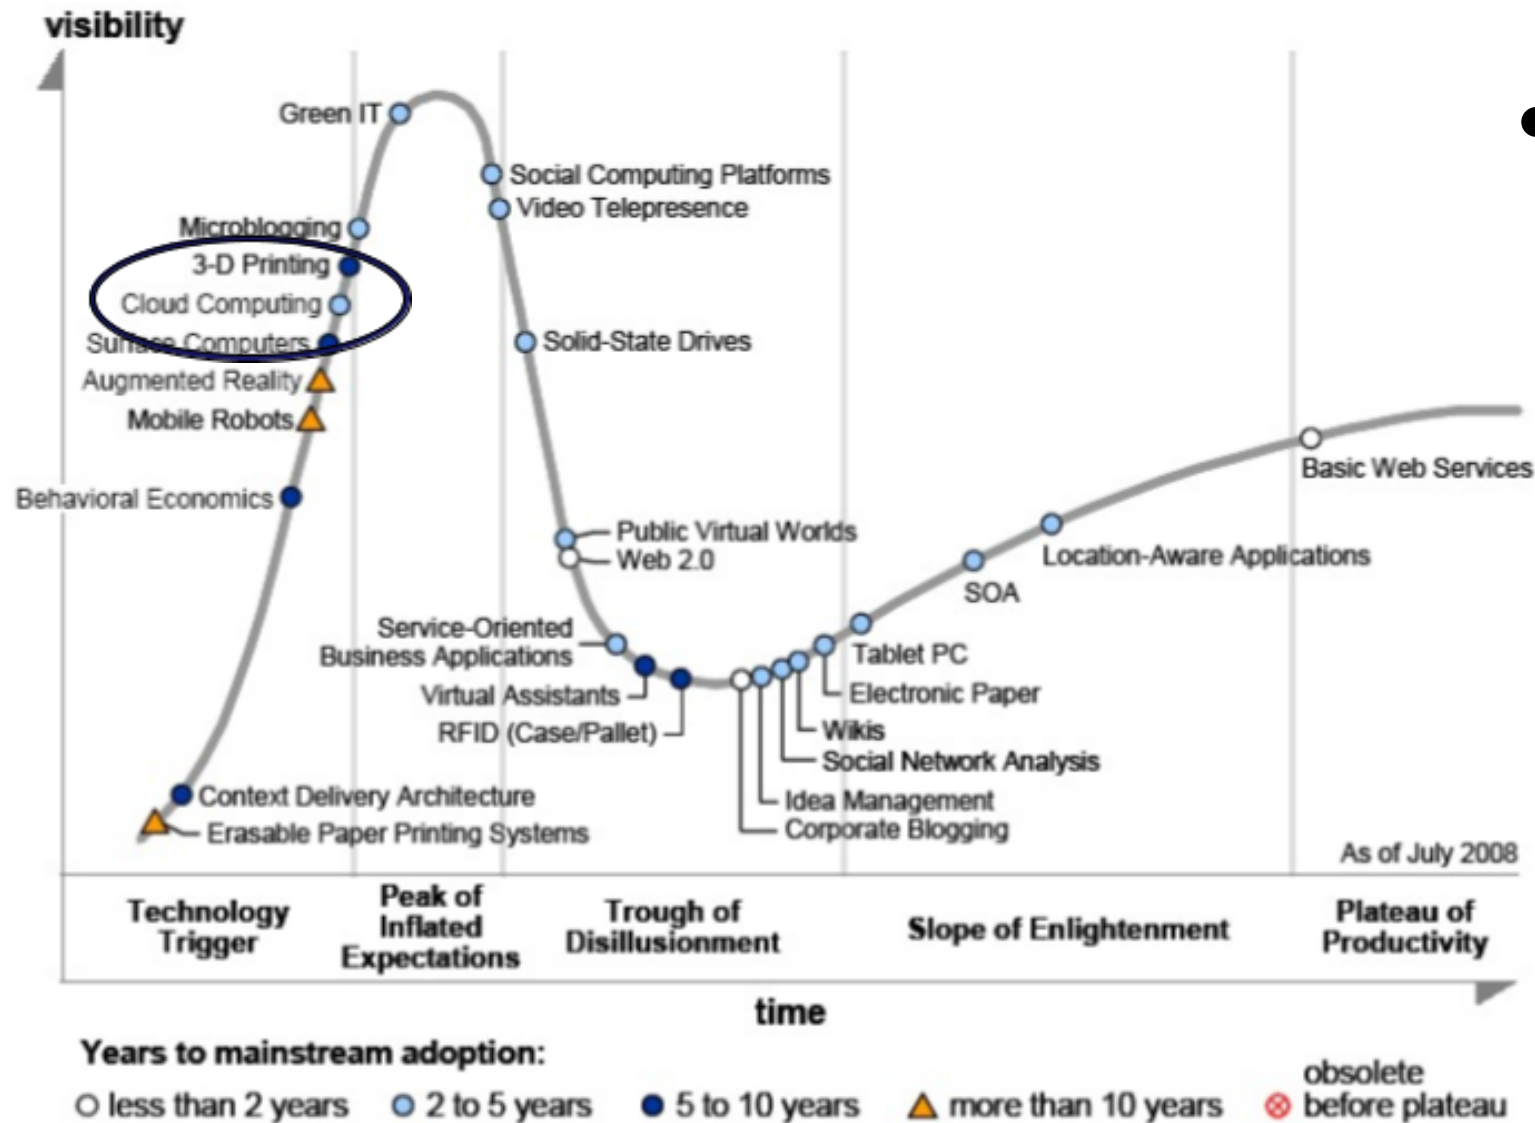
Search

## What is Cloud Computing




# On the Hype curve

Figure 1. Hype Cycle for Emerging Technologies, 2008



- Now probably at the top of the Hype – Oct 09

# Trendy...




- Source: <http://www.google.com/trends>



# Cloud formation



Amazon Elastic Compute Cloud (Amazon EC2) - Beta




- Slide adapted from Rich Wolski, UCSB

# White House is going to the Cloud



- Reduce costs ... See Apps.gov

HOME » SEARCH FOR  IN All Categories



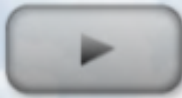
**Coming soon to Apps.gov**

You've got servers, developers, storage, testing, and upgrades to support your IT infrastructure. Need a better solution to reduce cost and speed implementation? Apps.gov Cloud IT Services can provide data storage, web hosting, and virtual machines all in the cloud - coming soon to Apps.gov!

**What is Cloud Computing?**

Want to learn more?

Watch this brief video for an overview of Cloud Computing to gain a better understanding of what it is and its benefits.




Watch the video now »

**What type of solution do you need?**

**Business Apps**

Your agency or service is complex and requires state-of-the-art software to get business done.


*GSA Cloud Business Apps has a solution!*



**Cloud IT Services**

Need a better solution to reduce cost and implement projects faster?

*GSA Cloud IT Services has the answer!*



# DOE (Magellan) and NASA too



**NEBULA Cloud Computing Platform** alpha

HOME ABOUT SERVICES BLOG

## Cloud Computing For a Universe of Data

A Cloud Computing environment combining a set of open-source components into a seamless self-service platform.

[More About NEBULA](#)

The screenshot shows the NEBULA Cloud Computing Platform interface. It features a header with the NASA logo and the text 'NEBULA Cloud Computing Platform alpha'. Below the header is a navigation bar with links for HOME, ABOUT, SERVICES, and BLOG. The main content area is titled 'Cloud Computing For a Universe of Data' and describes the platform as a seamless self-service environment. A 'More About NEBULA' button is present. To the right, there is a 'Bringing Science to a Web App Near You' section with a NASA logo and a description of the platform's capabilities. Below this is an 'Apps Spotlight' section featuring 'LCROSS' and 'TERRA'. The bottom of the banner has three icons: a lightning bolt for 'Rapid Deployment', a folder with an arrow for 'Scalable', and a padlock with a gear for 'Built for NASA'.

**Rapid Deployment**

**Scalable**

**Built for NASA**

# What is the Cloud ? The \*aaS



## ● **SaaS –Software as a Service –**

- Easy Access to hosted applications over the network. Most likely using your Browser
- API to these applications

## ● **PaaS –Platform as a Service –**

- Environment to deploy new applications
- Restricted capabilities offered
- API to this platform and access to SaaS API

## ● **IaaS –Infrastructure as a Service –**

- Access to Hardware resources
- API to make resource allocation requests

# References



## **"Above the clouds: A Berkeley view of cloud computing"**

<http://berkeleyclouds.blogspot.com/>

## **"A break in the clouds: towards a cloud definition"**

L.M Vaquero et al. SIGCOMM computer communication review, 2008.

<http://portal.acm.org/citation.cfm?id=1496100>

## **"An *EGEE* Comparative Study - Grid *cloud* comparative study"**

M-Elian Begin, 2009


# Key Features




- **You don't know what's behind but it works**
  - Transparency
- **You Pay what you use**
  - Utility pricing (real \$ ? )
- **You get what you ask for (On-demand)**
  - Instant provisioning of machines
- **It scales if you need more (Elasticity)**
  - How far does it scale ?
  - Doesn't this mean the underlying resources are underutilized ?
- **Virtualization is a key enabler of IaaS**




# Free resources or not ?



(a) Provisioning for peak load



(b) Underprovisioning 1



(c) Underprovisioning 2

Figure 2: (a) Even if peak load can be correctly anticipated, without elasticity we waste resources (shaded area) during nonpeak times. (b) Underprovisioning case 1: potential revenue from users not served (shaded area) is sacrificed. (c) Underprovisioning case 2: some users desert the site permanently after experiencing poor service; this attrition and possible negative press result in a permanent loss of a portion of the revenue stream.

# Deployment Models

## *Innovation in Cloud Computing Architectures*

| Model          | Definition                                                                                   | Examples of Deployment                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------|----------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Private</b> | Infrastructure is owned by a single organization and made available only to the organization | <ul style="list-style-type: none"> <li>• Optimize and simplify <b>internal operation</b></li> <li>• <b>SaaS/PaaS</b> support</li> <li>• IT consolidation within <b>large organizations</b> (Government Clouds, University Clouds...)</li> </ul>                                                                                                                                                                       |
| <b>Public</b>  | Infrastructure is owned by a single organization and made available to other organizations   | <ul style="list-style-type: none"> <li>• <b>Commercial cloud providers</b></li> <li>• <b>Community public clouds</b> by ICT service centers to enable scientific and educational projects to experiment with cloud computing</li> <li>• <b>Special purpose clouds</b> with dedicated capabilities (Science Clouds, HPC Clouds..)</li> <li>• <b>Regional clouds</b> to address regulatory or latency issues</li> </ul> |
| <b>Hybrid</b>  | Infrastructure is a composition of two or more clouds                                        | <ul style="list-style-type: none"> <li>• <b>Cloudbursting</b> to address peak demands</li> <li>• <b>Cloud Federation</b> to share infrastructure with partners</li> <li>• <b>Cloud Aggregation</b> to provide a larger resource infrastructure</li> </ul>                                                                                                                                                             |



# Main components/characteristics



## **Set of Hypervisors**

- Physical machines with a virtual machine monitor
- Xen or KVM ...or Hyper-V...or VMware ESx...

## **VM provisioning system**

- OpenNebula
- Nimbus
- Eucalyptus
- Platform ISF
- or even traditional schedulers like PBS/Maui.

## **Image distribution mechanism**

- Shared file system (e.g NFS, AFS, PVFS, Lustre...)
- Copy images (e.g scp, wget, Bittorrent)

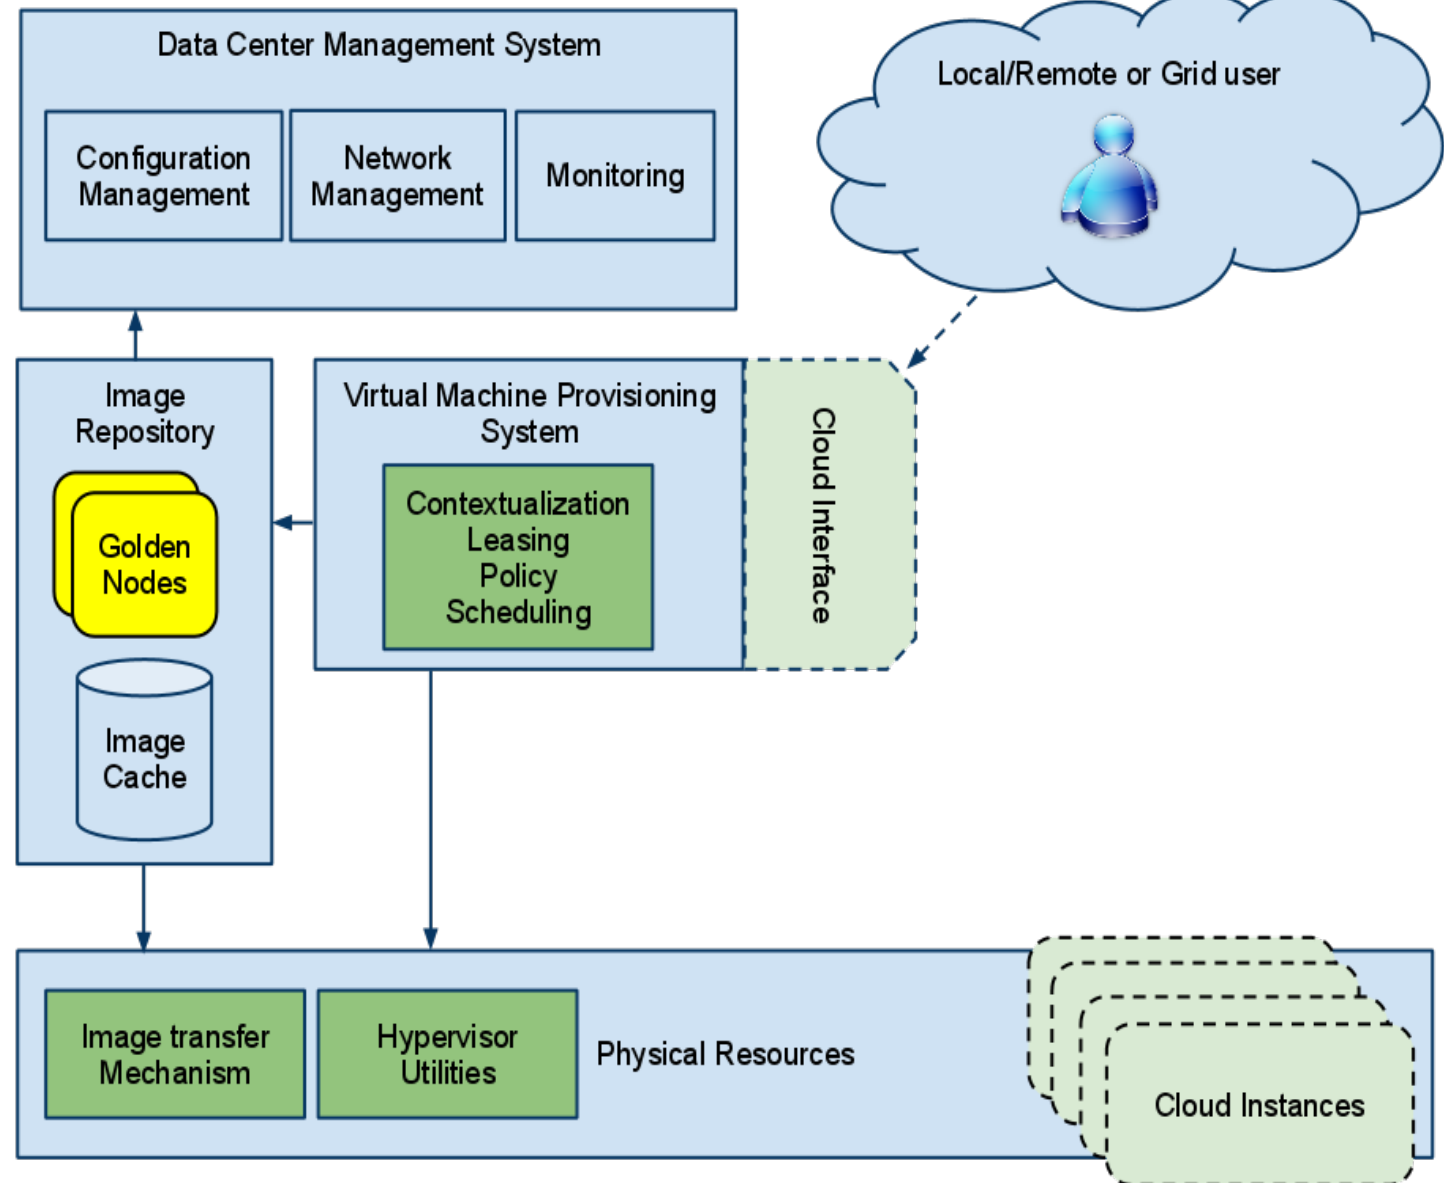
## **Networking**

- Private / Public bridged
- NAT

# CERN's Ixcloud architecture



- Image repository with Golden nodes.
- VM instances not quattor managed have finite lifetime
- Specific IP/MACs are pinned to hypervisors
- Currently testing two provisioning system: Opennebula and Platform ISF.



# Provisioning system



OpenNebula and Platform ISF are currently being evaluated.  
Results shown in this talk were obtained with OpenNebula.

OpenNebula out of the University Comptense of Madrid

- C/C++ core with Ruby drivers and command line interface
- Mysql and Sqlite backends
- Use ssh as communication between frontend and hosts
- XML-RPC API
- 
- Support for LVM contributed by CERN
- Enables *Hybrid clouds* (i.e instantiation on remote cloud providers)
- Implements subset of EC2 interface as well as upcoming OCCl interface for *Public cloud interface*

OpenNebula.org

# Building a Cloud: Interoperability Map

*Innovation in Cloud Computing Architectures*

Public



**OpenNebula**

OpenNebula



ElasticHosts



Private Cloud

Hybrid

# Comparison with Similar Technologies

*OpenNebula - Architecture, Current Status & Roadmap*

|                               | Platform ISF | VMware Vsphere | Eucalyptus | Nimbus    | OpenNebula                      |
|-------------------------------|--------------|----------------|------------|-----------|---------------------------------|
| Virtualization Management     | VMware, Xen  | VMware         | Xen, KVM   | Xen       | Xen, KVM, VMware                |
| Virtual Network Management    | Yes          | Yes            | No         | Yes       | Yes                             |
| Image Management              | Yes          | Yes            | Yes        | Yes       | Yes                             |
| Service Contextualization     | No           | No             | No         | Yes       | Yes                             |
| Scheduling                    | Yes          | Yes            | No         | No        | Yes                             |
| Administration Interface      | Yes          | Yes            | No         | No        | Yes                             |
| Hybrid Cloud Computing        | No           | No             | No         | No        | Yes                             |
| Cloud Interfaces              | No           | vCloud         | EC2        | WSRF, EC2 | EC2 Query<br>OGF OCCI<br>vCloud |
| Flexibility and Extensibility | Yes          | No             | Yes        | Yes       | Yes                             |
| Open Source                   | No           | No             | GPL        | Apache    | Apache                          |

# Cloud Interfaces...Standards ?



 [Sign in to the AWS Management Console](#)

[AWS](#)

[Products](#)

[Developers](#)

[Community](#)

[Support](#)

## Amazon Elastic Compute Cloud

[Home](#) > ... > [Amazon Elastic](#)

### Amazon EC2 API Tools

 [Printer Friendly](#)  [Save](#)



The API tools serve as a bridge between the Amazon EC2 API and the Open Cloud Computing Interface (OGFi) specification, allowing you to register and launch instances using the OGF API.

[Home](#) | [Roadmap](#) | [The Specification](#) | [Join](#)

**Submitted By:** [Dave](#)

### Download

[Download the Amazon](#)

See the related Amazon tools.

## OGF Open Cloud Computing Interface Working Group

### NEWS

**Have a look at our Open Cloud Computing Interface specification and provide us your comments:** [http://www.ogf.org/gf/docs/?public\\_comment](http://www.ogf.org/gf/docs/?public_comment)

### About the Open Cloud Computing Interface Working Group

Cloud computing currently is covered by three models offering Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS), which all involve the on-demand delivery of computing resources. There are a growing number of providers offering **IaaS solutions for elastic capacity**, whereby server "instances" are executed in their proprietary infrastructure and billed on a utility basis. The Open Cloud Computing Interface (OGFi) Working Group is a community of interest in the development of a standard interface for the delivery of cloud computing resources.

# econe-describe-images



```
[oneadmin@lxcloud bin]$ econe-describe-images -
H -K sebgoa -S '*****' -U http://lxcloud.
cern.ch:5555
```

| Owner  | ImageId      | Location           |
|--------|--------------|--------------------|
| sebgoa | ami-00000001 | /dev/xen_vg/cernvm |

# econe-run-instances



```
[oneadmin@lxcloud bin]$ econe-run-instances -H
-K sebgoa -S '*****' -U http://lxcloud.cern.
ch:5555 ami-00000001
```



# econe-describe-instances



```
[oneadmin@lxcloud bin]$ econe-describe-
instances -H -K sebgoa -S '*****' -U http:
//lxcloud.cern.ch:5555
```

# Issues to work on



- Leaving aside VM provenance/certification/distribution mechanisms and policies...
- VM accounting: who instantiated what VM for how long ?
- Pricing: what is the incentive to turn off your instance ?
- Authentication to "EC2" interface with grid proxies ?
- Additional service to maintain in CERN-IT and other sites...

### European Projects on Cloud Computing Infrastructures



RESERVOIR

EU grant agreement 215605  
**Service and Sw Architectures  
and Infrastructures**  
(2008-2011)

#### **Resources and Services Virtualization without Barriers**

- Open source technology to enable deployment and management of complex IT services across different administrative domains



StratusLab

Proposal in negotiation  
**e-Infrastructure**  
(2010-2012)

#### **Enhancing Grid Infrastructures with Cloud Computing**

- Simplify and optimize its use and operation, providing a more flexible, dynamic computing environment for scientists.
- Enhance existing computing infrastructures with “IaaS” paradigms



BonFire

Proposal in negotiation  
**New Infrastructure Paradigms  
and Experimental Facilities**  
(2010-2013)

#### **Building Service Testbeds on FIRE**

- Design, build and operate a multi-site cloud-based facility to support research across applications, services and systems targeting services research community on Future Internet

# Conclusions



Ixcloud is a reality but not yet a production service

Production processes are being worked on

- See Ulrich's talk

Policies are being worked on

- See Tony's talk

- Working on running CERNVM as a "standalone" appliance
- Working on running virtual batch worker nodes.