

HP Labs Seminar
Palo Alto, CA
November 2nd, 2010

OpenNebula

*Leading Innovation in
Cloud Computing Management*


Ignacio M. Llorente and Rubén S. Montero

DSA-Research.org
Distributed Systems Architecture Research Group
Universidad Complutense de Madrid

Acknowledgments




The research leading to these results has received funding from the European Union's Seventh Framework Programme ([FP7/2007-2013]) under grant agreement n° 215605 (RESERVOIR Project)



Private Cloud Computing => A “Public Cloud behind the firewall”


- Simplify and optimize internal operations
- Service flexibility and elasticity
- Higher utilization & operational savings
- Security concerns

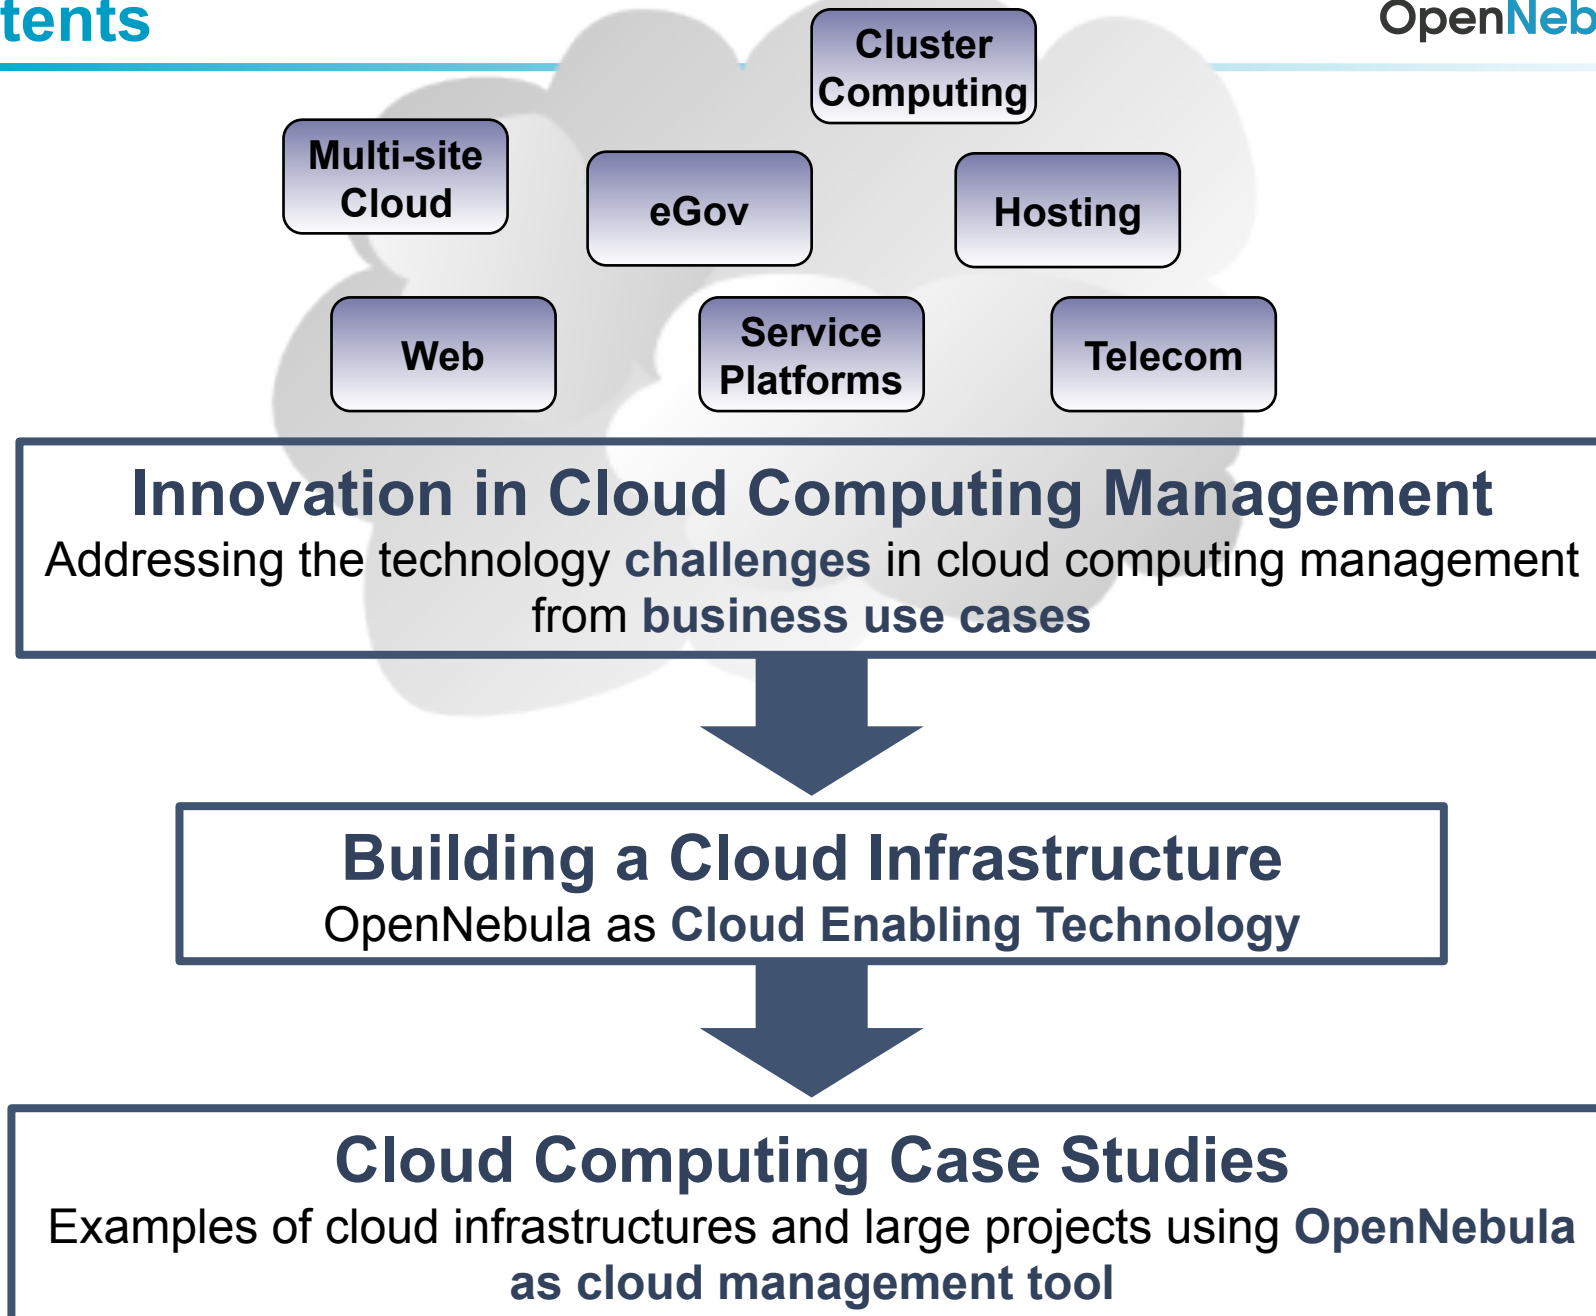


Hybrid Cloud Computing => Utility Computing dream made a reality!

- Supplement the capacity of the Private Cloud


Cloud Manager to Orchestrate the Complexity of a Datacenter






Profile of Service Workloads

- Multi-tier service as **basic management entity**



- Compute, storage and network **capacity**
- **Level of coupling** between service instances
- **Security and placement constraints**
- **Automatic configuration** of service instances



Service as Groups of VMs


- Service **components** in VMs
- **Inter-connection** relationship
- **Placement constraints**

Cloud Interface

- Management of **VM images, VM instances, and virtual networks**
- Support for **standard and common interfaces** (OGF OCCl, Amazon EC2 and VMware vCloud)

Comprehensive Management of the Cloud

- **Administration interface** for the centralized management of the cloud
 - Physical infrastructure with hosts and clusters management
 - Users, and authorization and authentication
 - VM images, VM instances and virtual networks
- Definition of workload and resource-aware **allocation policies** such as energy efficiency, load balancing, affinity-aware, capacity reservation
- **Secure multi-tenancy** and isolation
- Site policy enforcement with user **quota management**
- **Accounting** to “charge” users based on usage or to guarantee fair share of resources among users
- **Highly reliable, efficient and scalable back-end**

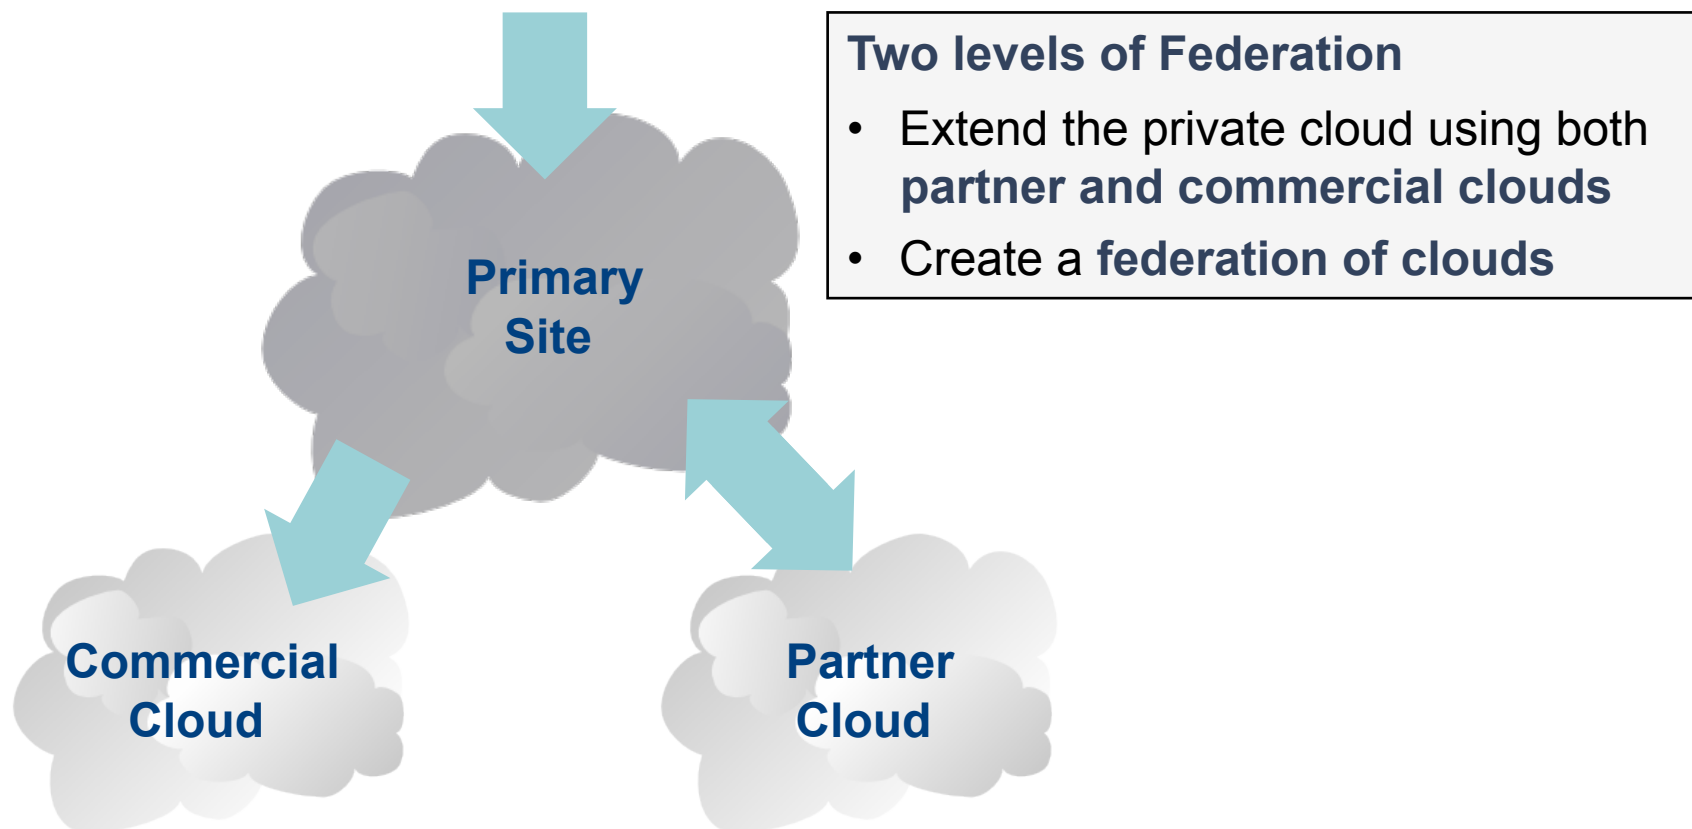


Scalable back-end

- Virtualization
- Storage
- Networking

Hybrid Cloud Computing and Federation


- **Cloudbursting** at infrastructure layer, fully transparent to users
- **Scale-out decisions** are taken by infrastructure administrators according to business policies



- Cloud Computing is an **evolution of existing data centers**
- One solution can not fit all data-center and user requirements and constraints

**Constraints from
Existing Infrastructure
and Processes**

**Requirements from
Usage and Deployment
Scenarios**



“One solution does not fit all requirements and constraints. There cannot be turnkey quick cloud solutions”

Open Architecture, Interfaces and Code

- **Integration** with existing **processes** and **management tools** in the data center
- Support any **security, monitoring, storage, networking and virtualization infrastructure service**
- **Integration with any product and service** in the virtualization/cloud ecosystem such as cloud providers, hypervisors, virtual image managers, service managers, management tools, schedulers...
- Support to **build any type of deployment**: private, public, hybrid and community clouds
- **Easy to extend** to support new functionality and **to embed** into other Cloud applications and platforms
- Based on **standards** to avoid vendor lock-in and to enable interoperability
- **Truly open-source**, not open core
- **Liberal open-source license**

Flagship International Projects in Cloud Computing

Result of many years of research and development in efficient and scalable management of virtual machines on large-scale distributed infrastructures.



Open-source Toolkit

Open platform for innovation to research the challenges that arise in **enterprise cloud management**, and **production-ready** tool in both academia and industry

- **Started in 2005, first release in march 2008, and ONE 2.0 just released**
- **Open-source** released under Apache v2.0, packaged for main Linux distributions
- Mailing lists for **best-effort support** and **open development framework**
- Development and roadmap definition **driven by the community and projects**
- Active and engaged **open community and ecosystem**
- **> 3,000 downloads/month** (not including code repository and Ubuntu)
- Used in many **production environments**, distributed in **commercial solutions** and availability of **commercial professional support by C12G Labs**
- **Long-term sustainability** ensured by project funding and commercial sponsors

Capabilities for Cloud Management

Most advanced open-source toolkit offering unique features to administer the complexity of large-scale distributed infrastructures

Capabilities for Integration

Open, flexible and extensible architecture, interfaces and components that fit into any existing data center

Capabilities for Production Environments

Scalability and performance tested on very large-scale infrastructures consisting of thousands of cores, with the security and fault tolerance levels required in production

Leverage the Vibrant Cloud Ecosystems

Leverage the ecosystems being built around OpenNebula and the most common cloud interfaces, Amazon AWS, OGC OCCl and VMware vCloud

Fully Open Source Cloud Software

OpenNebula is NOT a feature or performance limited edition of an Enterprise version.
OpenNebula is truly open, and not open core.

Openness

- Open architectures
- Open interfaces
- Open code

Adaptability

- Modular architectures

The logo for OpenNebula.org is centered on the slide. It features a stylized, multi-colored cloud shape in shades of blue and grey. The text "OpenNebula.org" is overlaid on the cloud, with "Open" in black, "Nebula" in blue, and ".org" in black.

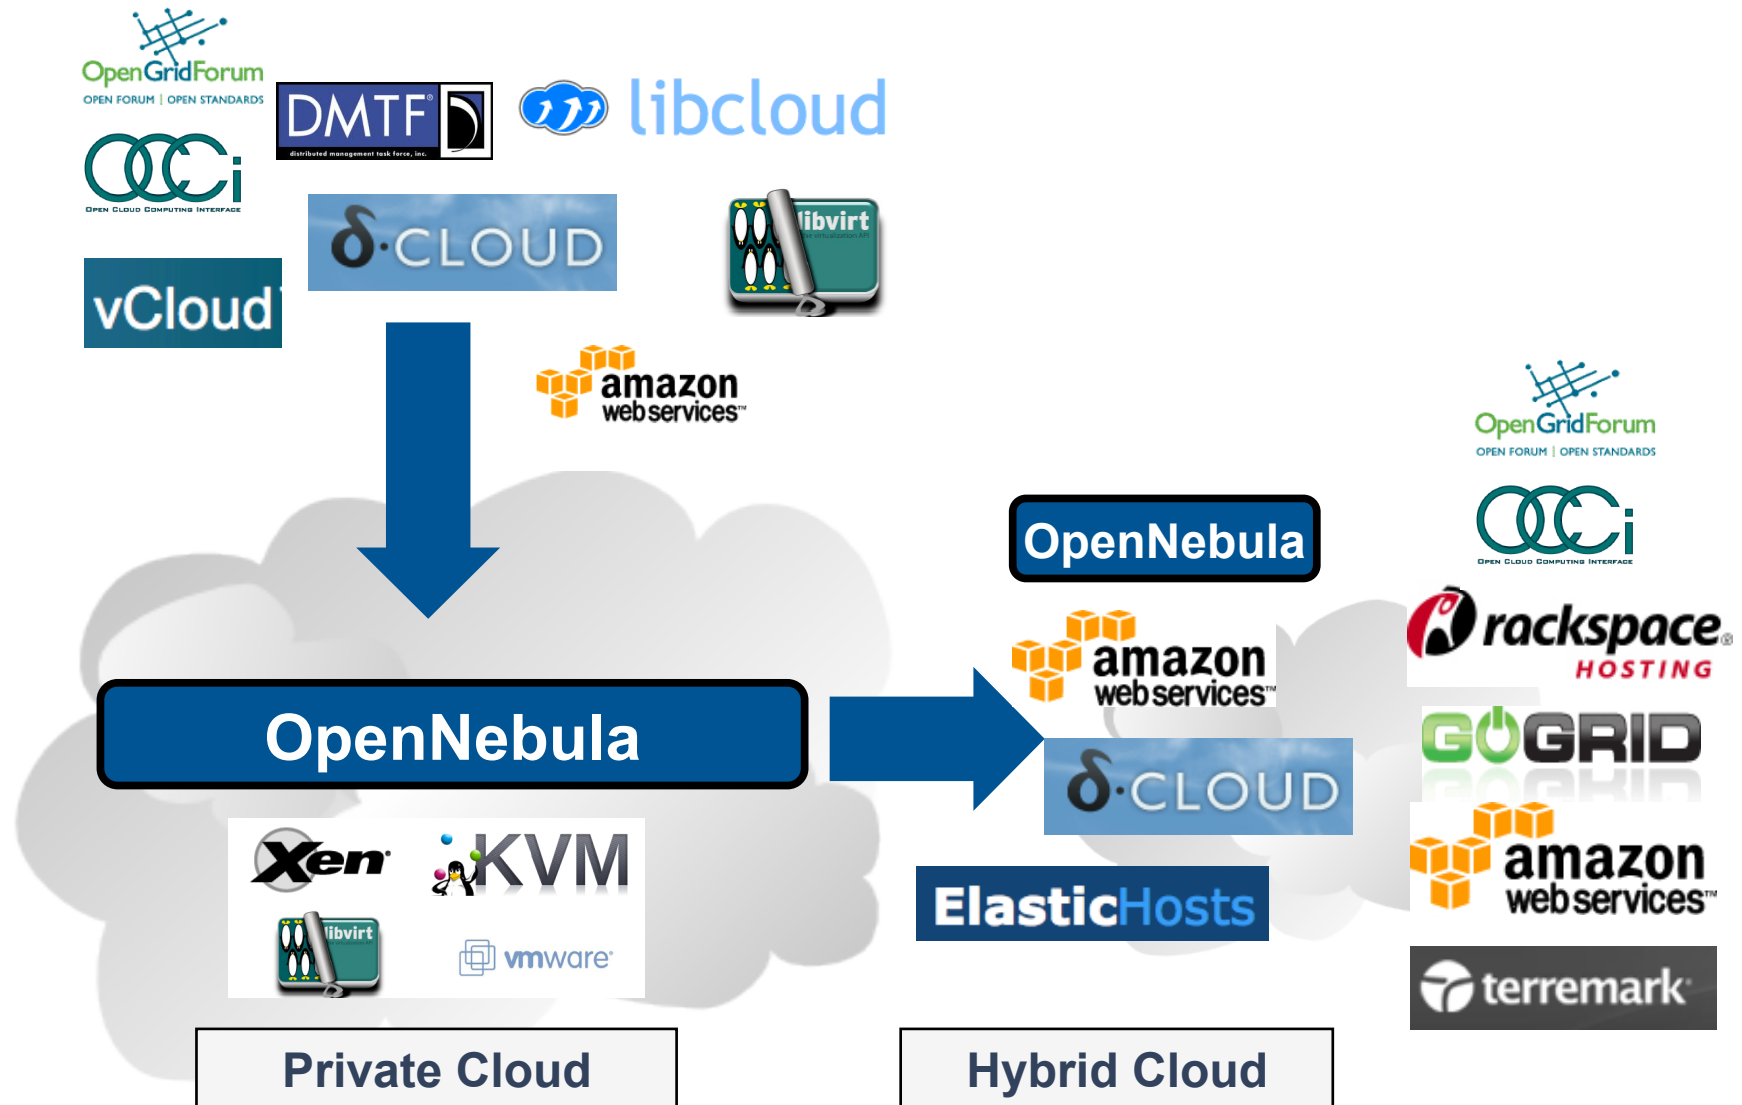
OpenNebula.org

Standardization

- Use standards
- Implement standards


Portability

- It can be installed in any hardware and software




A Highly Modular Architecture to Fit into any Existing Datacenter


- Cloud Computing is an **evolution of existing data centers**
- One solution can not fit all data-center and user requirements and constraints
- Open, flexible and extensible architecture
- Provide basic components, but allow them to be easily replaceable by others



Examples of Components in the Ecosystem



Adopt Standards



Open Source Community

- Open architecture and interfaces
- Open code and liberal license
- Open community and ecosystem

OpenNebula.org

Management Tool



Innovation Tool



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

Private Cloud to Support Grid Site



- **Goal:** Execution of virtualized Grid sites in D-Grid and EGEE
- **Details:** The D-Grid Resource Center Ruhr (DGRZR) runs an OpenNebula private cloud on 248 blades and 1,984 cores with Xen


Public HPC Cloud



- **Goal:** OpenNebula is used to support the execution of virtual clusters and HPC applications
- **Details:** SARA High Performance Computing Center uses OpenNebula in its new HPC Cloud service



- **Goal:** Easier management and new computing models in the batch farm
- **Example of Integration with Existing Infrastructure Environment**
 - **Configuration Management:** Quattor with lifecycle management and “self-notification” in OpenNebula
 - **Network Management:** Adapted to address network infrastructure requirements regarding fixed IP/MAC leases in each box
 - **Storage Management:** New LVM transfer scripts and a very fast parallel scp to push images to all the hosts





Source: CERN IT-PES/PS Group: Sebastien Goasguen, Ulrich Schwickerath, Ewan Roche and Belmiro Moreira

- **Goal:** Meet the growing demands for high performance, low cost, high scalability, high reliability of China Mobile IT Infrastructure (computing, storage); and the demands of China Mobile to deliver Internet business and services



- **Details:** 4,096 cores, Xen, Ganglia, and Hadoop




Source: China Mobile's Presentation at OpenCirrus Meeting



Agreement 215605 (2008-2011)
Service and Sw Architectures
and Infrastructures

Resources and Services Virtualization without Barriers

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



StratusLab


EU grant agreement RI-261552
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; and enhance existing computing infrastructures with “IaaS” paradigms



www.StratusLab.eu






Agreement 257386 (2010-2013)
**New Infrastructure Paradigms
and Experimental Facilities**

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



www.BonFIRE-Project.eu



Source: BonFIRE Project



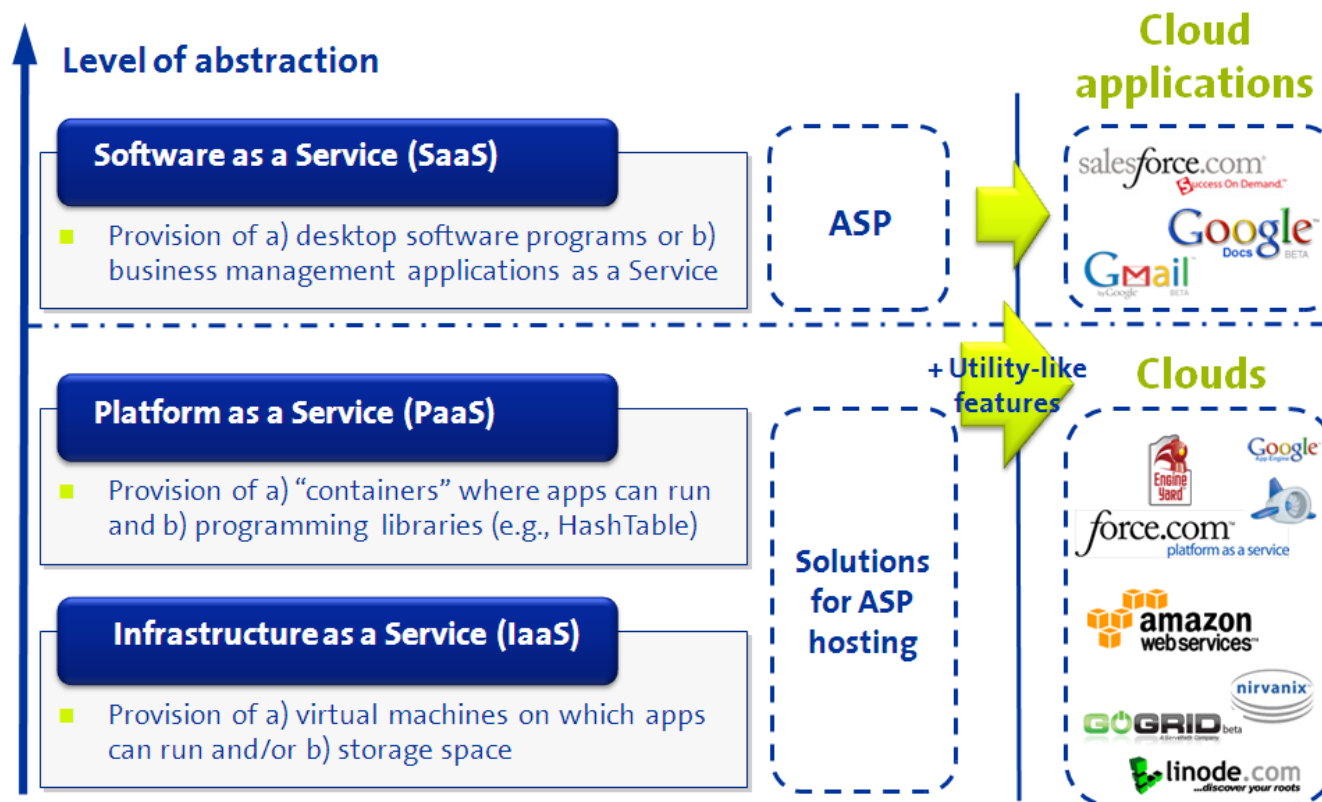
Agreement 258862 (2010-2013)
**Service and Sw Architectures
 and Infrastructures**

Building the PaaS Cloud of the Future

- Create an advanced PaaS Cloud platform which supports the optimized and elastic hosting of Internet-scale multi-tier applications



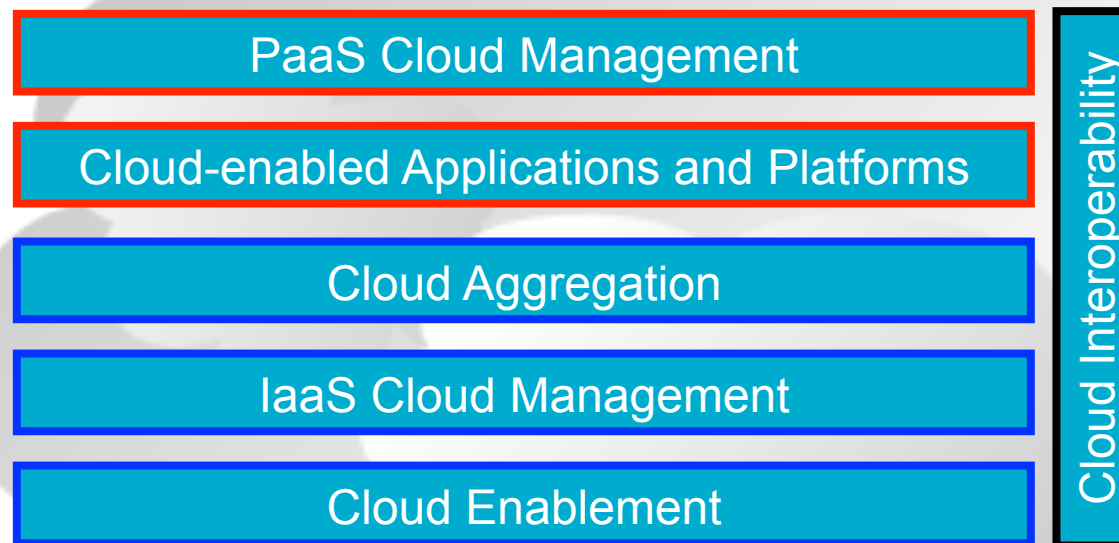
4caast.morfeo-project.org



Source: 4CaaS Project

IT Resources will be the Next Utility

- **Future enterprise datacenters will operate like hybrid Clouds** combining local with public Cloud-based infrastructure to enable highly scalable hosting environments
- **Public Clouds will be supported by a network of geographically distributed datacenters** for high availability, end-user service proximity, legal and policy issues...
- **Growing number of domain specific and regional Cloud providers** implementing a utility computing business model



Use the Technology and Give us Feedback

- Support through several mailing lists or describe your use case in our blog
- Report bugs and make feature requests

Spread our Word

- Spread the word about OpenNebula and open source cloud computing

Contribute to the Development

- Open development infrastructure
- Provide patches for bug fixes or enhancements

Contribute to the Quickly Growing Ecosystem

- Submit a new tool or extension to the OpenNebula ecosystem

Sponsor the Community

- Provide funds or resources to support development or to organize events

Innovation Alliance

- Collaboration in open-source and innovation in cloud computing management

More info, downloads, mailing lists at

The screenshot shows the OpenNebula.org website. At the top, it says "OpenNebula.org" and "The Open Source Toolkit for Cloud Computing". Below this is a navigation menu with links for Home, About, Documentation, Software, Support, Community, Cloud, Dev, and Blog. The main content area features a "The Truly Open-Source, Leading and Most Advanced Cloud Software" section with a list of features: Private cloud with Xen, KVM and VMware; Hybrid cloud with Amazon EC2, and other providers through Deltacloud (from ecosystem); and Public cloud supporting EC2 Query, OGF OCCT and vCloud (from ecosystem) APIs, and much more. There is also a "Getting Started" section with three steps: 1. Download OpenNebula, 2. Read the Documentation, 3. Engage the Community. A "Try it now!" button is visible in the "OpenNebula 2.0 RC1" section.

The screenshot shows the RESERVOIR website. At the top, it says "RESERVOIR Resources and Services Virtualization without Barriers". Below this is a navigation menu with links for Home, What is Reservoir?, Technical Info, Downloads, Training, Demos & Videos, Media Centre, Events & Presentations, Blogs, and News. The main content area features a "RESERVOIR: Business Driven Research" section with a sub-heading "How the Research Community is facilitating on-demand services for business". Below this is a "Business Driven Research" section with a sub-heading "Technical Information" and a "RESERVOIR Framework Downloads" section with a sub-heading "RESERVOIR Demos & Videos".

Research References

- B. Rochwerger, J. Caceres, R.S. Montero, D. Breitgand, E. Elmroth, A. Galis, E. Levy, I.M. Llorente, K. Nagin, Y. Wolfsthal, "The RESERVOIR Model and Architecture for Open Federated Cloud Computing", **IBM Systems Journal**, Vol. 53, No. 4. (2009)
- B. Sotomayor, R. S. Montero, I. M. Llorente and I. Foster, "Virtual Infrastructure Management in Private and Hybrid Clouds", **IEEE Internet Computing**, September/October 2009 (vol. 13 no. 5)



The research leading to these results has received funding from the European Union's Seventh Framework Programme ([FP7/2007-2013]) under grant agreement n° 215605 (RESERVOIR Project)