

**OSDC 2011**  
Nuremberg, Germany  
5-7 April, 2010

# OpenNebula Tutorial

**Constantino Vázquez Blanco**

DSA-Research.org

Distributed Systems Architecture Research Group  
Universidad Complutense de Madrid


## Acknowledgments



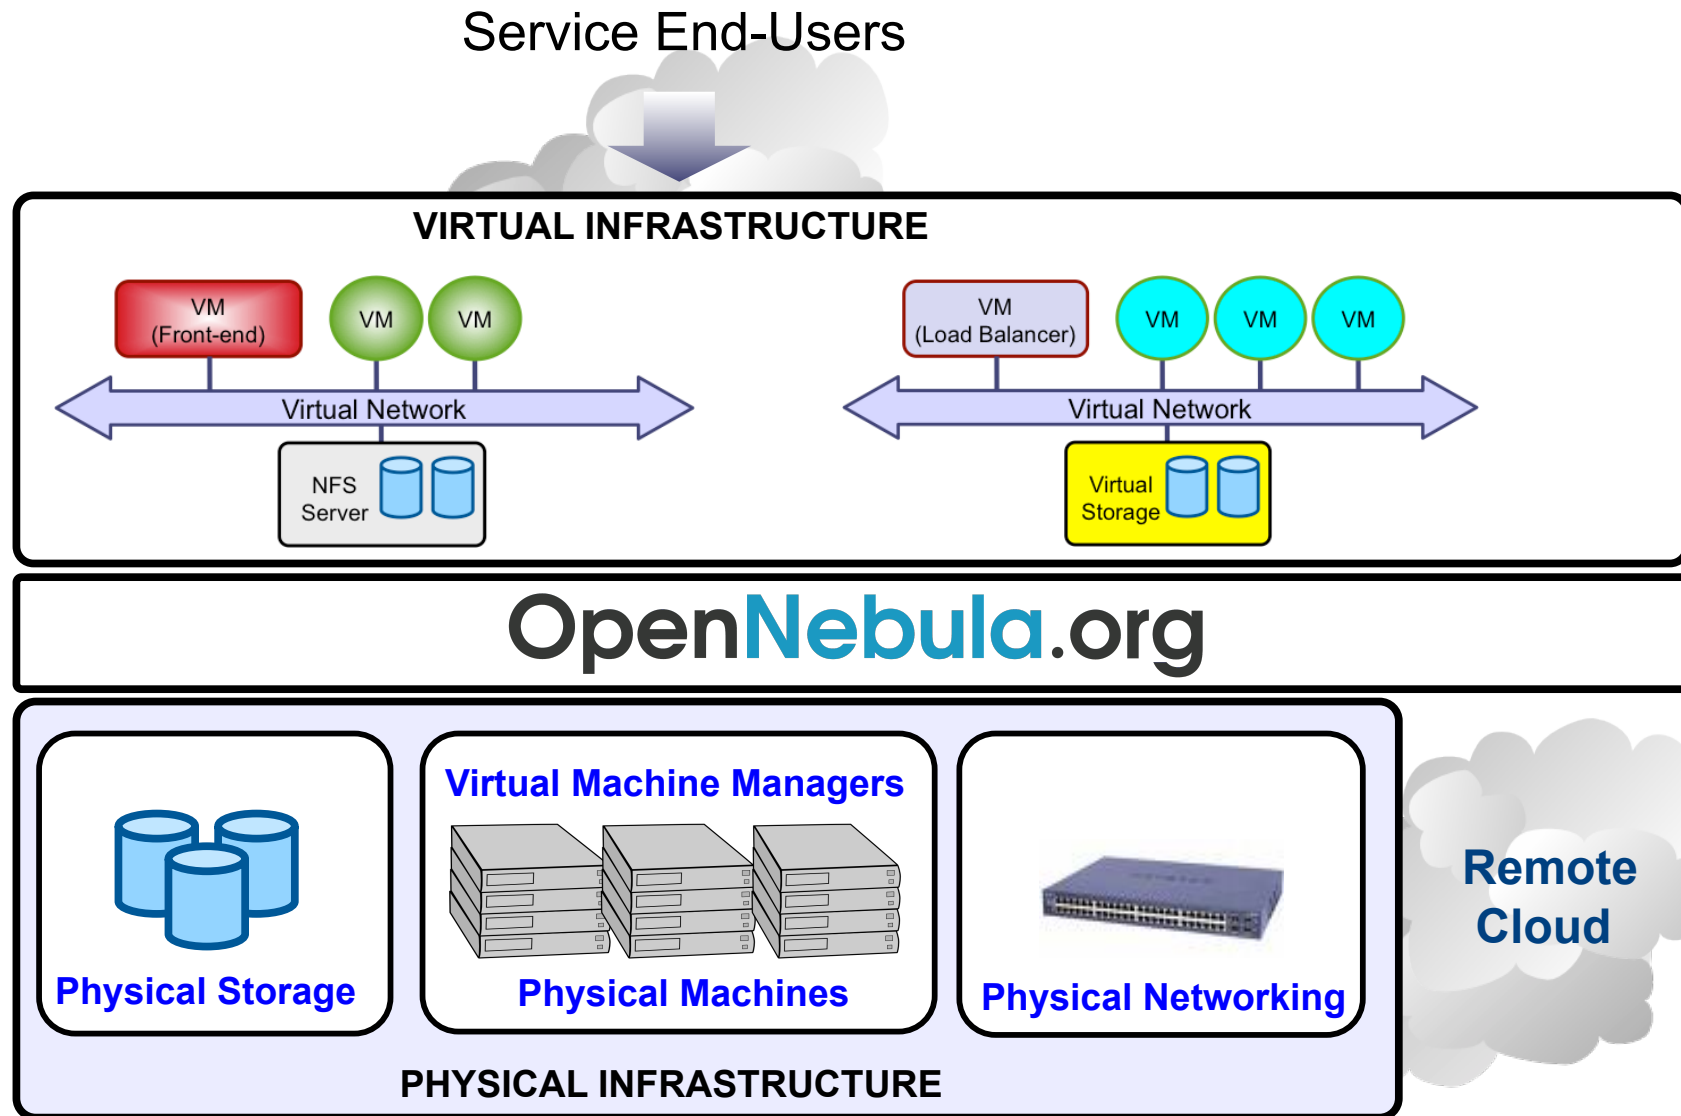
*The research leading to these results has received funding from the European Union's Seventh Framework Programme ([FP7/2007-2013] ) under grant agreement n° 215605 (RESERVOIR Project)*

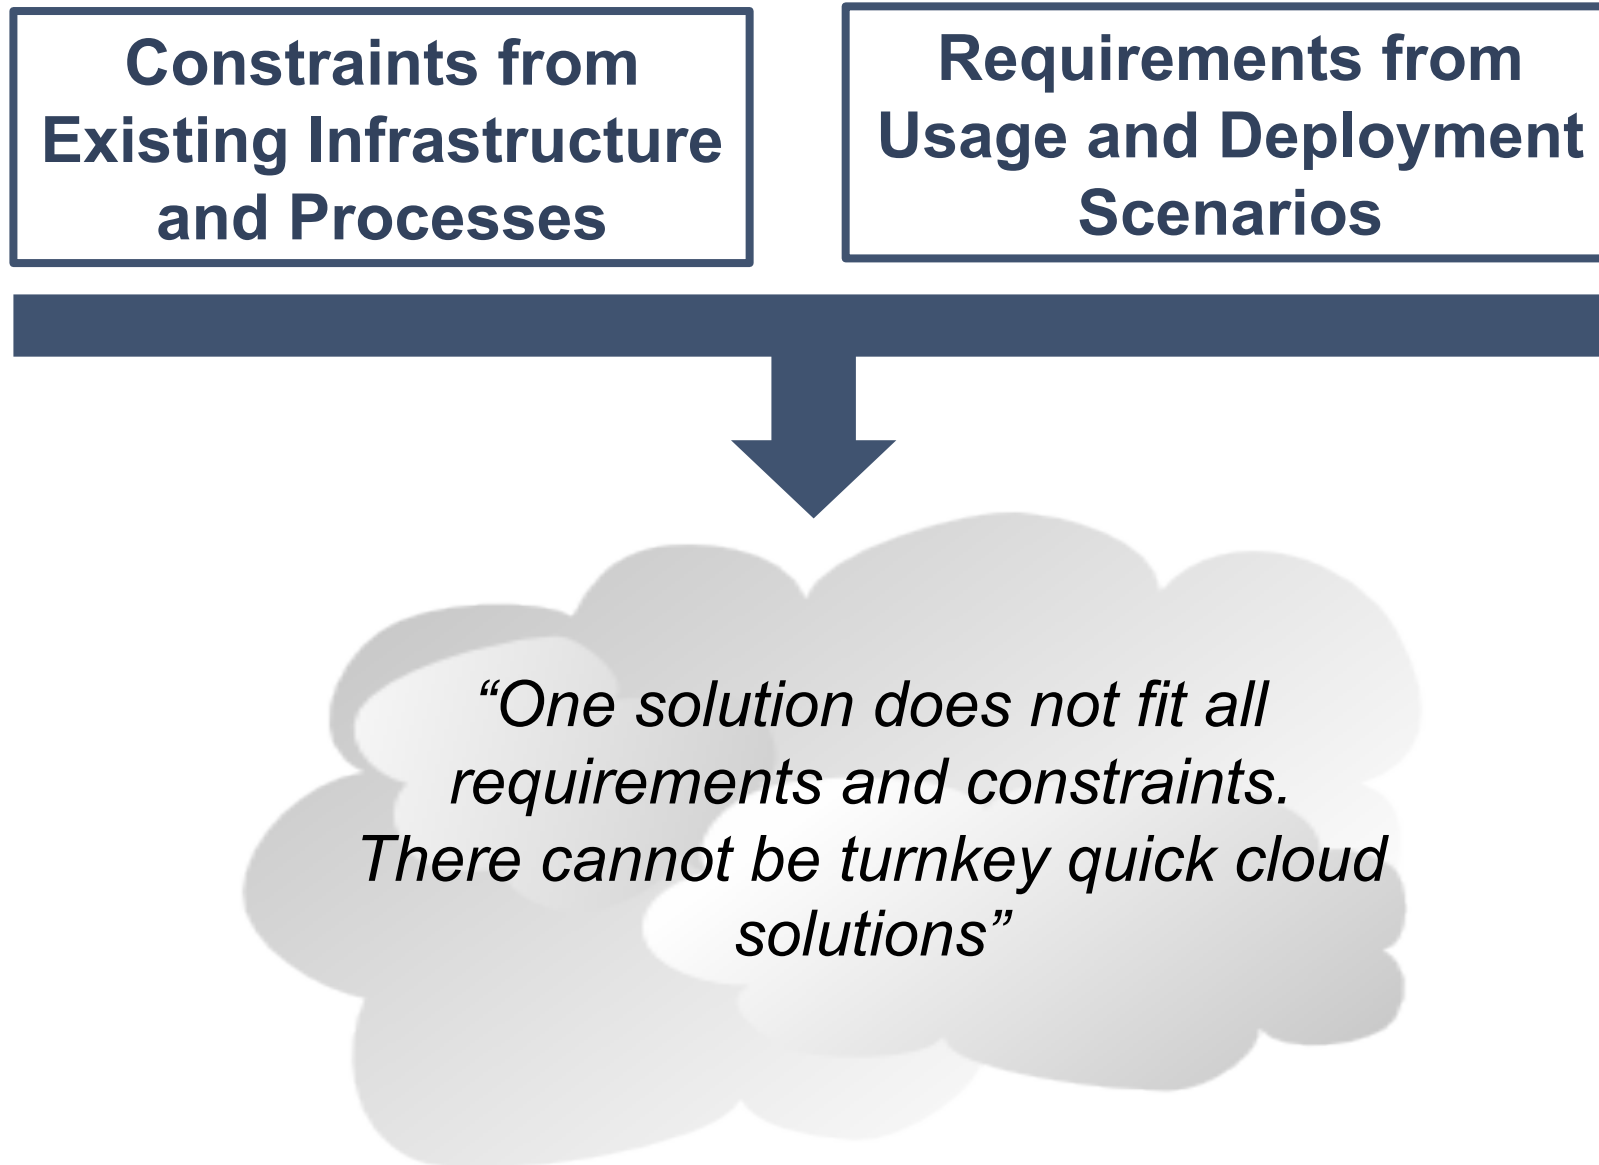
- **OpenNebula Overview**
- **Installation and Configuration**
- **Private Cloud: Administration and Basic Usage**
- **Hybrid Cloud: Overview**
- **Public Cloud: Overview**
- **New Features in v2.2**

# OpenNebula Overview



## Cloud Manager to Orchestrate the Complexity of a Datacenter





## 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 cloud management, and production-ready tool in both academia and industry

- **Started in 2005, first release in march 2008, and ONE 2.2 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.



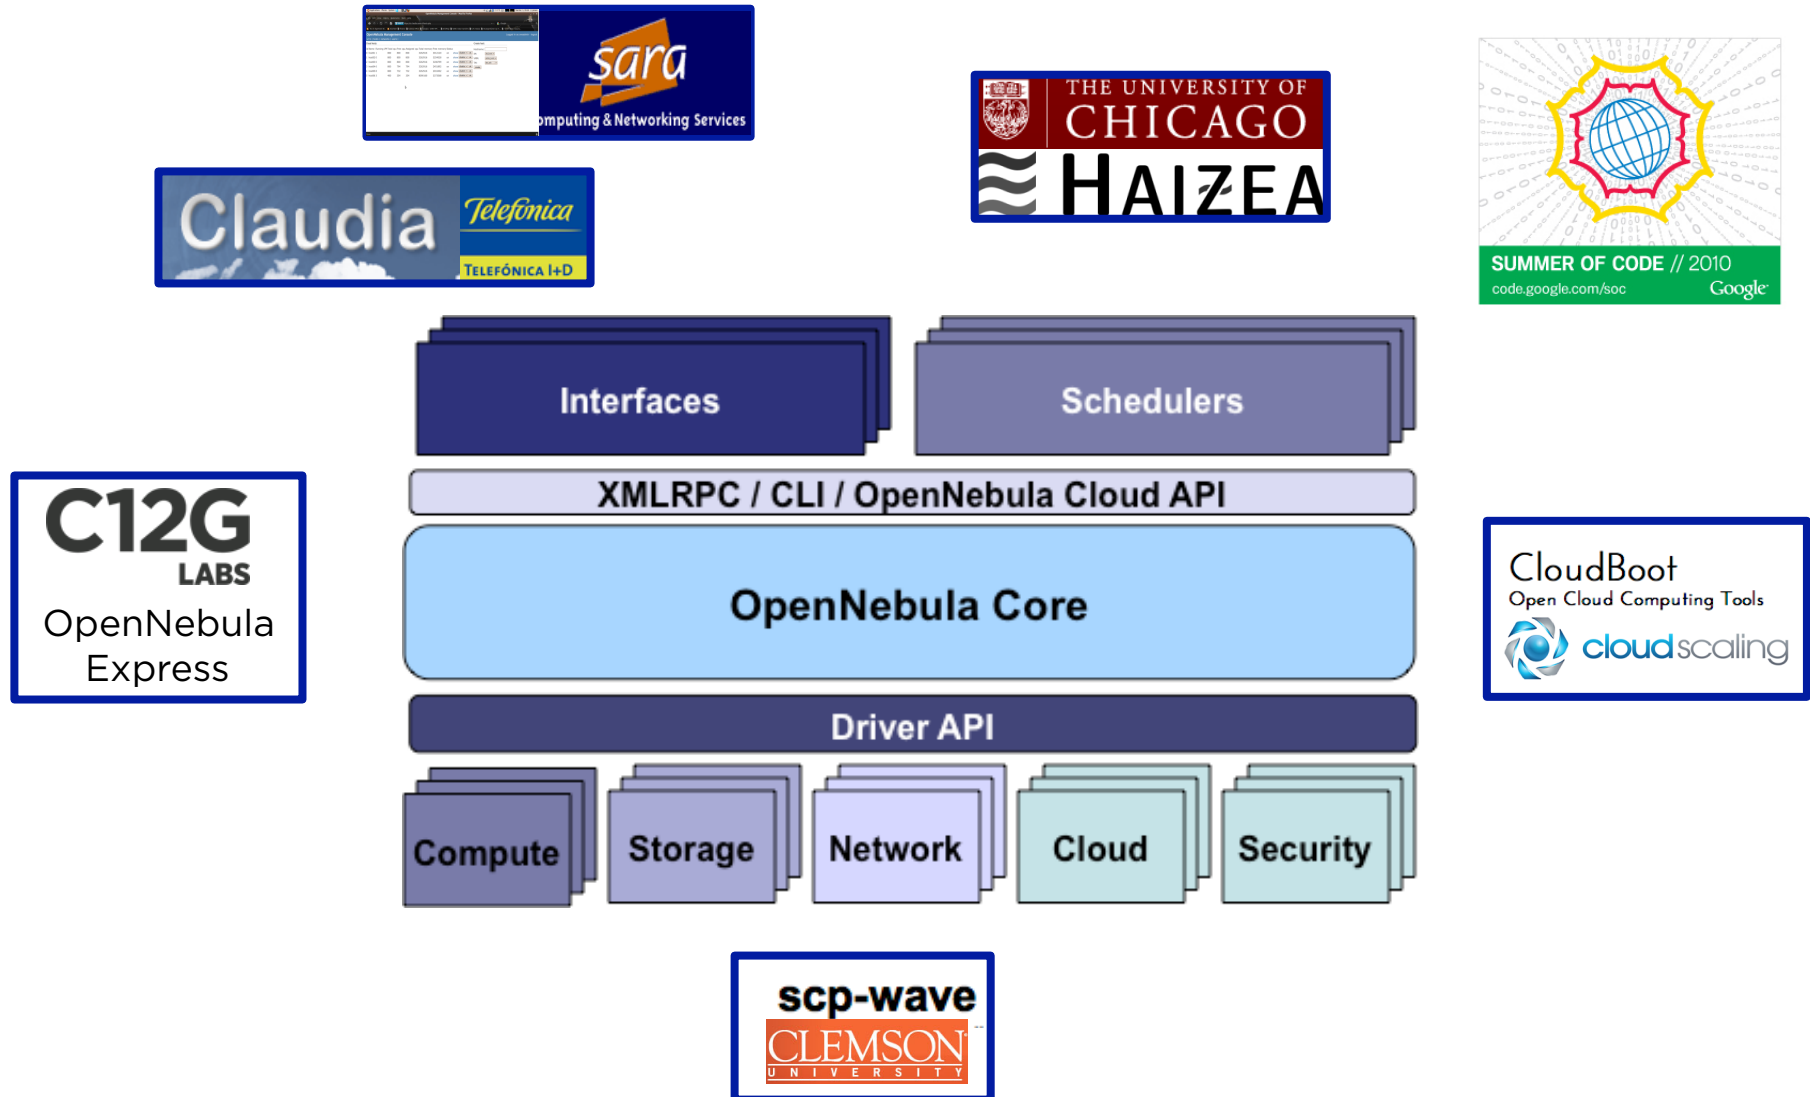
## A Highly Modular Architecture to Fit into any Existing Datacenter

- 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



# Installation and Configuration





# Private Cloud Administration and Basic Usage

- Add physical **Hosts**
- Create **Users**
- Define VM **Images**
- Define **Virtual Networks**
- Create **Virtual Machine** Templates
- Launch **Virtual Machines**



- A Host in OpenNebula
  - Is a resource with capacity (CPU, memory)
  - Is able to run VMs
  - Must be configured as an OpenNebula node
- Host definition
  - **Name** of the physical host
  - **Drivers**
    - **Information**, to extract resource usage
    - **Virtualization**, to interact with the hypervisor
    - **Transfer**, to perform image provision
- From OpenNebula 2.0 onwards – Logic clusters
  - Groups hosts to enforce placement policies
- Hosts are managed with the `onehost` command

- A User in OpenNebula
  - Is an entity that owns resources
- User definition
  - **Name**
  - **Password**
- From OpenNebula 2.0 onwards – Auth module
  - Pluggable architecture
  - Drivers for ssh and ldap
- Users are managed with the `oneuser` command

- An Image in OpenNebula (> 2.0)
  - Is a VM image contained in a Repository
  - Has metadata on how to use the VM image
- Image definition
  - **Name**, of the image
  - **Source**, of the file
  - Can be **persistent** or **public**
- Images are managed with the `oneimage` command

- A Virtual Network in OpenNebula
  - Defines a separated MAC/IP address space to be used by VMs
  - Each virtual network is associated with a physical network through a bridge
  - Virtual Networks can be isolated (at layer 2 level) with ebtables and hooks
- Virtual Network definition
  - **Name**, of the network
  - **Type**
    - **Fixed**, a set of IP/MAC leases
    - **Ranged**, defines a network range
  - **Bridge**, name of the physical bridge in the physical host where the VM should connect its network interface.
- Virtual Networks are managed with the `onevnet` command

- A Virtual Machine in OpenNebula
  - A **capacity** in terms memory and CPU
  - A set of **NICs** attached to one or more virtual networks
  - A set of **disk images**, to be *transferred* to/from the execution host.
  - A **state file** (optional) or recovery file, with the memory image of a running VM plus some hypervisor specific information.
- Virtual Machines are defined in a **VM template** (text file)
- All the files (logs, images, state files...) are stored in `$ONE_LOCATION/var/<VM_ID>`
- Virtual Machines are managed with the `onevm` command

```
NAME      = ttylinux
CPU       = 0.1
MEMORY   = 64

DISK      = [ image = "base ttylinux" ]

NIC       = [ NETWORK = "One-TD" ]

#This may be useful to debug your VMs (can use also console)

GRAPHICS = [
  type = "vnc",
  listen = "localhost",
  port = "5902",
  keymap="es" ]
```

- Using a Virtual Network with your VMs
  - Define NICs attached to a given virtual network. The VM will get a NIC with a free MAC in the network and attached to the corresponding bridge

```
#A VM with two interfaces each one in a different vlan
```

```
NIC=[NETWORK="One-TD"]
```

```
NIC=[NETWORK="One-TD-Invisible"]
```

```
#Ask for a specific IP/MAC of the Red vlan
```

```
NIC=[NETWORK="One-TD", IP=192.168.$CN.140]
```

- Prepare the VM to use the IP. Sample scripts to set the IP based on the MAC are provided for several Linux distributions.

## IP-MAC address correspondence

IP: 10.0.1.2


MAC: 02:01:0A:00:01:02

oned.conf      IP Address

# Hybrid Cloud Overview



- VMs can be local or remote
- VM connectivity has to be configured, usually VPNs



- External Clouds are like any other host
- Placement constraints

```
IM_MAD = [  
  name          = "im_ec2",  
  executable    = "one_im_ec2",  
  arguments     = "im_ec2/im_ec2.conf" ]  
                # No. of instances of each type
```

```
VM_MAD = [  
  name          = "vmm_ec2",  
  executable    = "one_vmm_ec2",  
  arguments     = "vmm_ec2/vmm_ec2.conf",  
                # Defaults, e.g. keypair  
  type         = "xml" ]
```

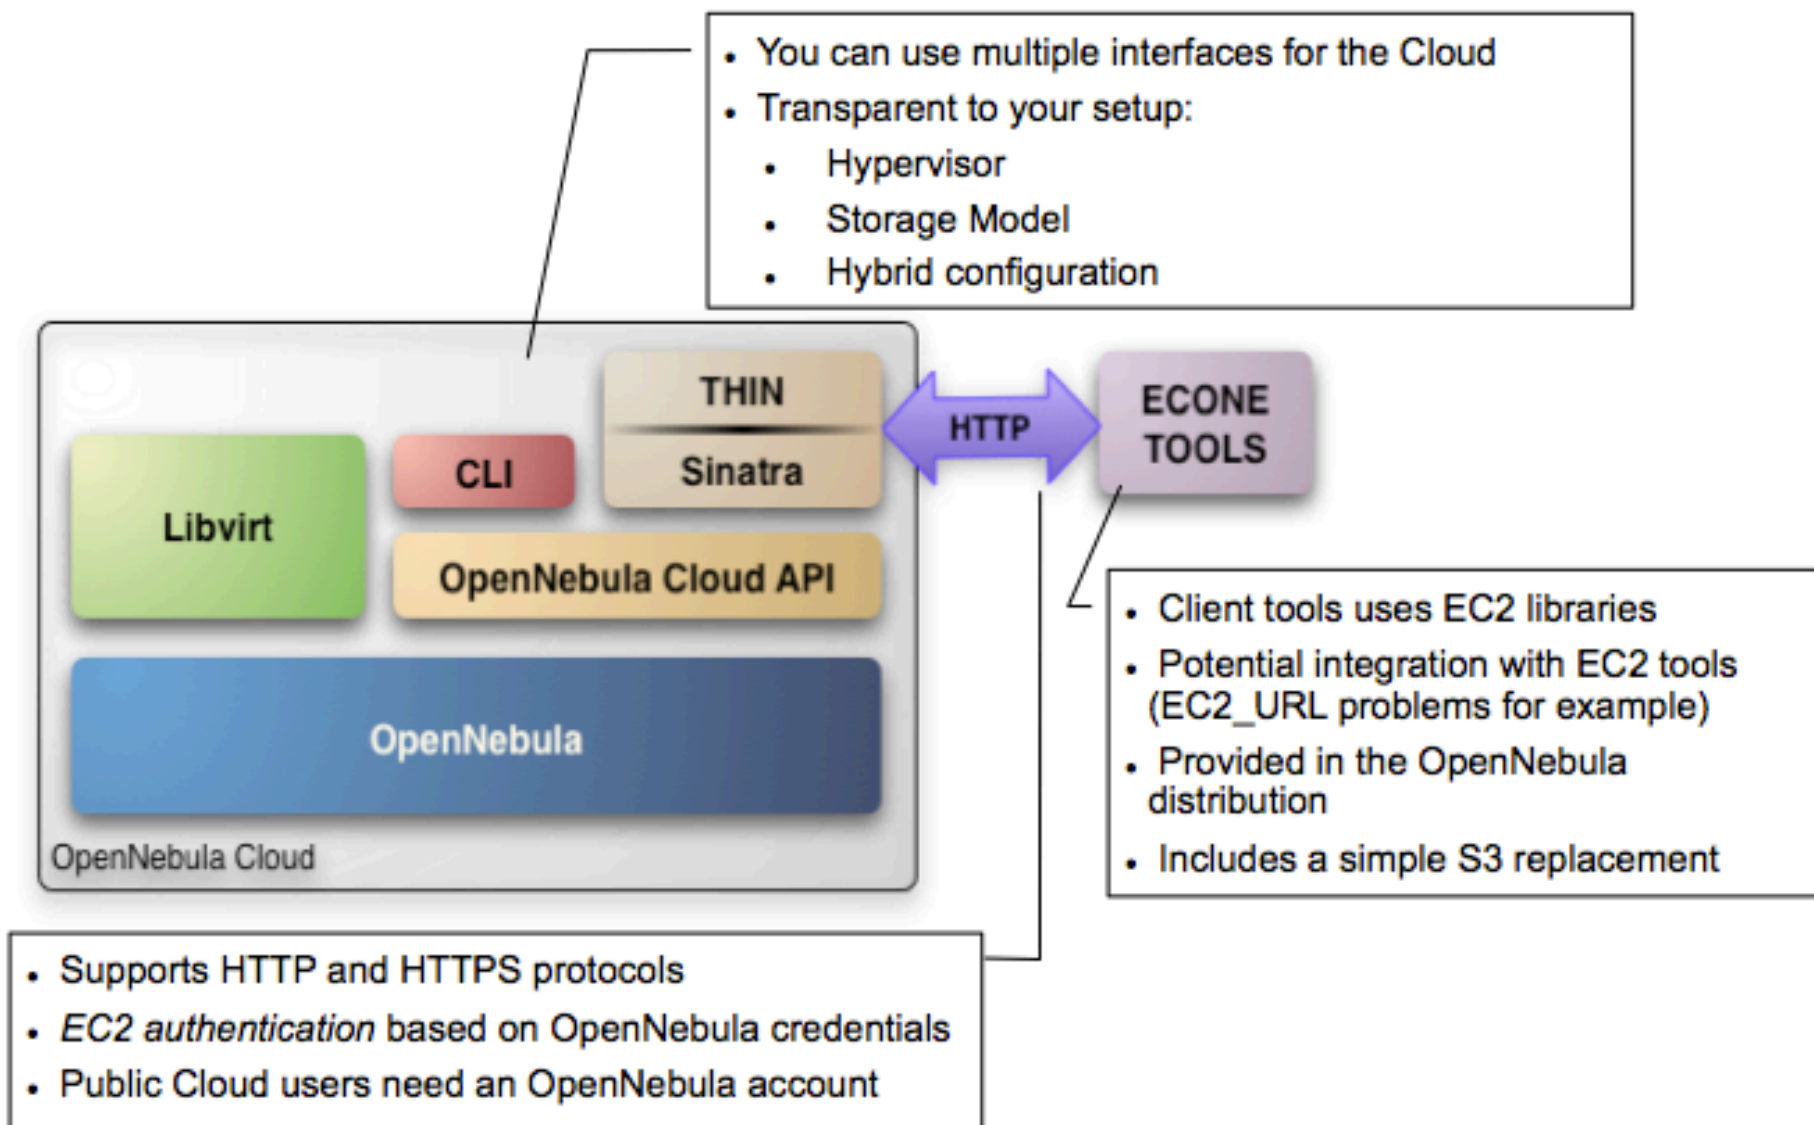
```
TM_MAD = [  
  name          = "tm_dummy",  
  executable    = "one_tm",  
  arguments     = "tm_dummy/tm_dummy.conf" ]
```

*# No actual transfers are made by OpenNebula to EC2*

- Virtual Machines can be instantiated locally or in EC2
  - The template must provide a description for both instantiation methods.
  - The EC2 counterpart of your VM (`AMI_ID`) must be available for the driver account
  - The EC2 VM template attribute:

```
EC2 = [  
  AMI           = "ami_id for this VM",  
  KEYPAIR       = "the keypair to use the instance",  
  AUTHORIZED_PORTS = "ports to access the instance",  
  INSTANCETYPE  = "m1.small...",  
  ELASTICIP     = "the elastic ip for this instance",  
  CLOUD         = "host (EC2 cloud) for this VM"  
]
```

# Public Cloud Overview



```
$ more $ONE_LOCATION/etc/econe.conf
# OpenNebula administrator user, the one_auth contents
USER=oneadmin
PASSWORD=onecloud

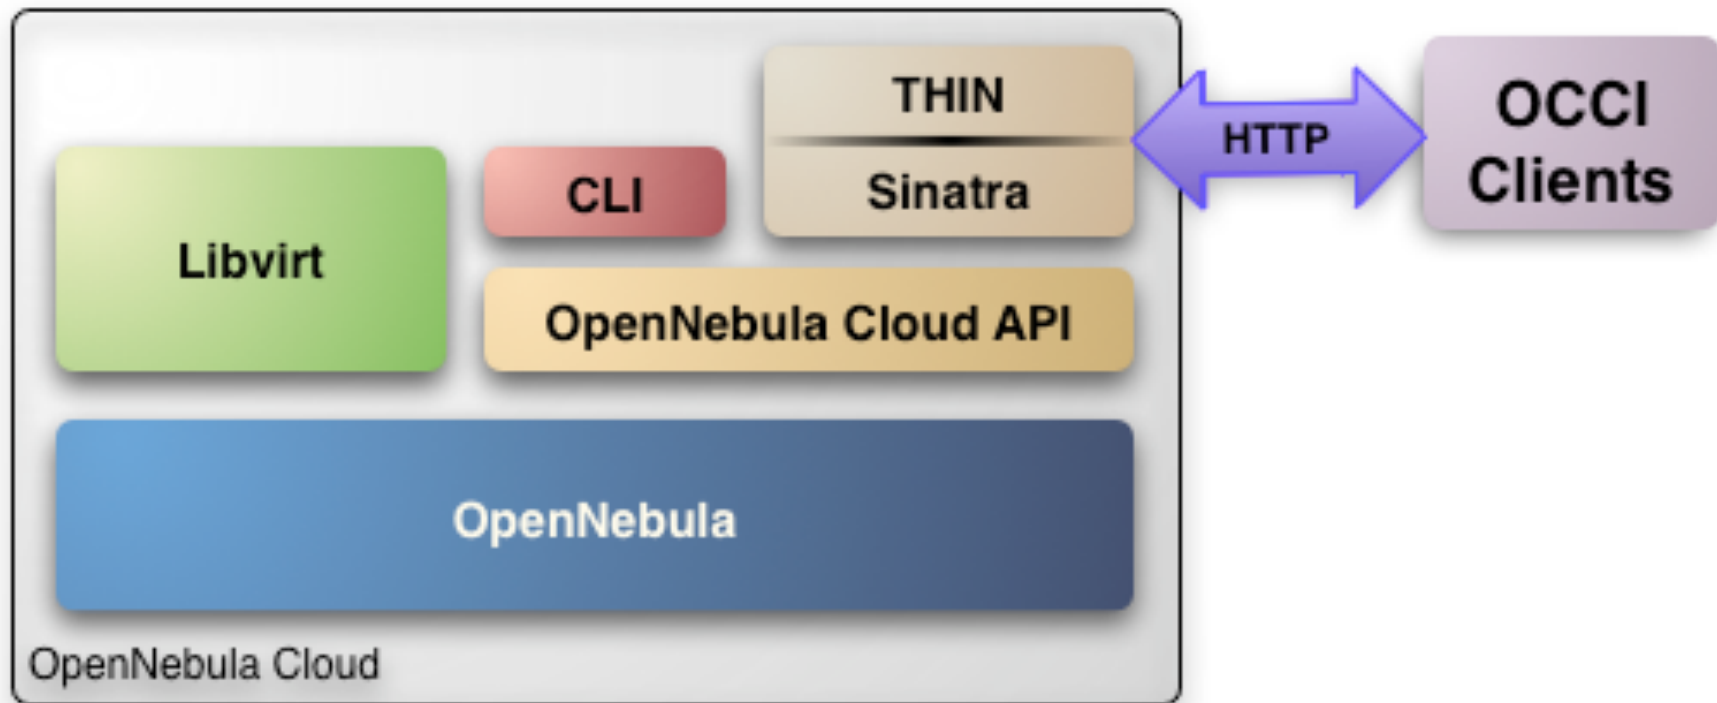
# OpenNebula sever contact information
ONE_XMLRPC=http://localhost:2633/RPC2

# Host and port where econe server will run keep FQDNs
SERVER=node-y.opennebula.org
PORT=4567

# Configuration for the image repository
# IMAGE_DIR will store the Cloud images, check space!
DATABASE=/srv/cloud/one/var/econe.db
IMAGE_DIR=/srv/cloud/public_repo/

# VM types allowed and its template file
VM_TYPE=[NAME=m1.small, TEMPLATE=m1.small.erb]
```

- The **econe-tools** are a subset of the functionality provided by the `onevm` utility, and resembles the `ec2-*` cli
- **Image** related commands are:
  - `econe-upload`, place an image in the Cloud repo and returns ID
  - `econe-describe-images`, lists the images
  - `econe-register`, register an image not really needed in 1.4
- **Instance** related commands are:
  - `econe-run-instances`, starts a VM using an image ID
  - `econe-describe-instances`, lists the VMs
  - `econe-terminate-instances`, shutdowns a VM
- User authentication is based in the OpenNebula credentials
  - `AWSAccessKeyId` is OpenNebula's username
  - `AWSSecretAccessKey` is OpenNebula's password





- Similar configuration as ECONE
- **REST protocol**
  - **Storage**
    - Upload: using a multi-part HTTP POST.
    - Retrieve: using a HTTP GET.
  - **Network**
    - Upload: using a HTTP POST.
    - Retrieve: using a HTTP GET.
  - **Compute**
    - Upload: using a HTTP POST.
    - Update: using a HTTP PUT.
    - Retrieve: using a HTTP GET.

# New Features in v2.2

- New GUI

The screenshot displays the OpenNebula Sunstone dashboard. At the top, the header includes 'OpenNebula Sunstone' on the left, and 'Documentation | Support | Community' and 'Welcome oneadmin | Sign Out' on the right. A left sidebar contains navigation links: 'Dashboard' (highlighted), 'Hosts & Clusters', 'Virtual Machines', 'Virtual Networks', 'Images', and 'Users'. The main content area features six summary cards, each with a title and a '+' icon in the top right corner. The 'Hosts' card shows 'Total: 5' and 'Active: 0'. The 'Clusters' card shows 'Total: 2'. The 'Virtual Machines' card shows 'Total: 4', 'Running: 2', and 'Failed: 1'. The 'Virtual Networks' card shows 'Total: 2' and 'Public: 0'. The 'Images' card shows 'Total: 1' and 'Public: 1'. The 'Users' card shows 'Total: 4'. At the bottom of the dashboard, a footer contains the text: 'Copyright 2002-2011 © OpenNebula Project Leads (OpenNebula.org). All Rights Reserved.'

- New Ganglia probes
  - Monitorization of physical and VMs through ganglia
  - Improves scalability
- Image Catalog exposed through OCCl
- New Fault Tolerance mechanisms
  - Physical Hosts
  - VMs
  - OpenNebula Service

## 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 section titled "The Truly Open-Source, Leading and Most Advanced Cloud Software" with a list of features: "Fully open source (not open core), flexible, extensible and with excellent performance and scalability to manage tens of thousands of VMs:", "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." To the right of this text is a cloud diagram with labels for Scalability, Openness, Robustness, Security, Community, and Flexibility. Below the features is a "Getting Started" section with three numbered steps: "1 Download OpenNebula", "2 Read the Documentation", and "3 Engage the Community". At the bottom of this section is a button that says "Your Cloud on Ubuntu and RHEL/CentOS in Few Minutes". On the right side of the page, there is a section for "OpenNebula 2.0 RC1" with a "Try it now!" button, a "Featured Quotes" section with a quote from Sebastian Goasguen, and an "Announcements" section listing recent updates.

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 section titled "RESERVOIR: Business Driven Research" with a sub-heading "How the Research Community is facilitating on-demand services for business". Below this is a paragraph describing the project: "RESERVOIR is a European Union FP7 funded project that will enable massive scale deployment and management of complex IT services across different administrative domains, IT platforms and geographies by providing a foundation for a service-based online economy, using virtualization technologies, transparently provisioned and managed on an on-demand basis at competitive costs with high quality of service. [Read More]". Below the text are four columns of content: "Business Driven Research" with a photo of people in a meeting, "Technical Information" with a network diagram, "RESERVOIR Framework Downloads" with a download icon, and "RESERVOIR Demos & Videos" with a play button icon.

## 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)

# Use Cases

## Adopt Standards



## Open Source Community

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

OpenNebula.org

## Management Tool



中国移动通信  
CHINA MOBILE

## Innovation Tool



| Model          | Definition   | Cloud Cases   |
|----------------|--|---|
| <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>Science 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 (HPC Clouds..)</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>   |



## Private Cloud to Support Grid Site



- **Goal:** Execution of a virtualized Grid site 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




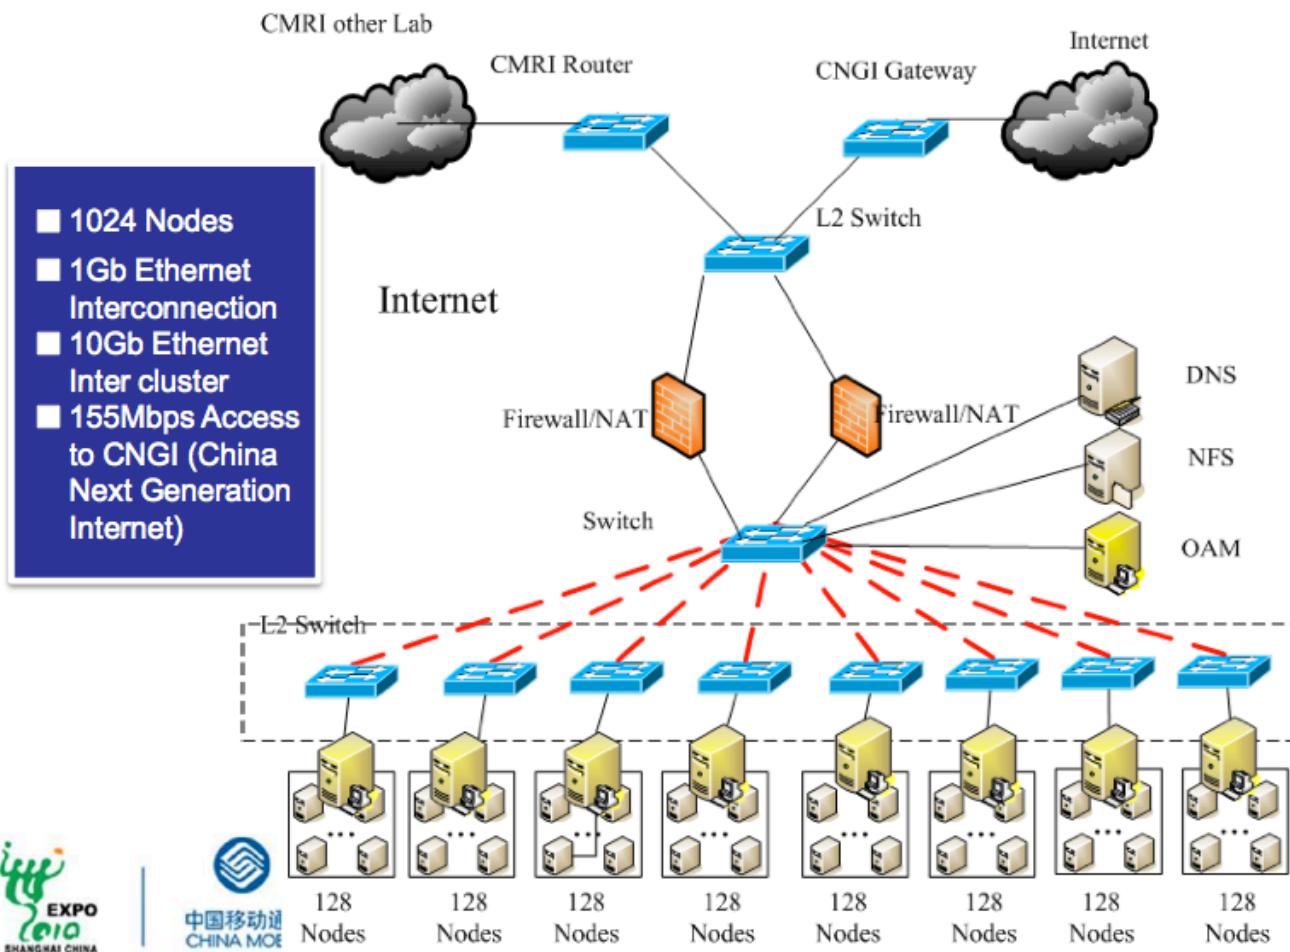
Source: CERN IT-PES/PS Group: Sebastien Goasguen, Ulrich Schwickerath, Ewan Roche and Belmiro Moreira



- **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



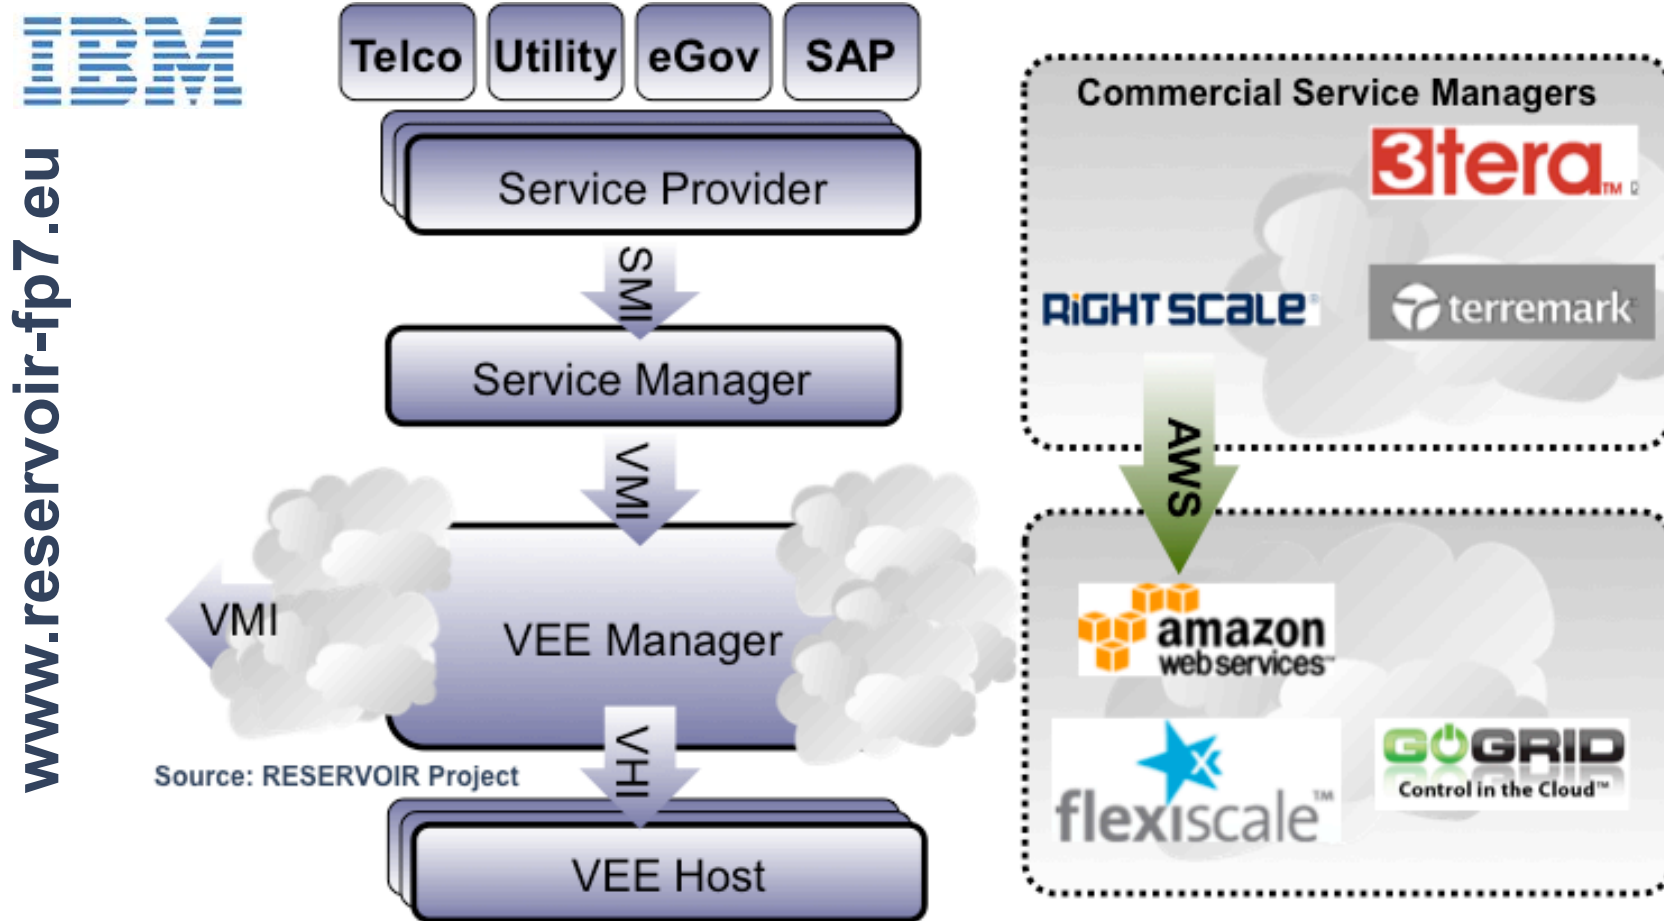
Source: China Mobile's Presentation at OpenCirrus Meeting



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





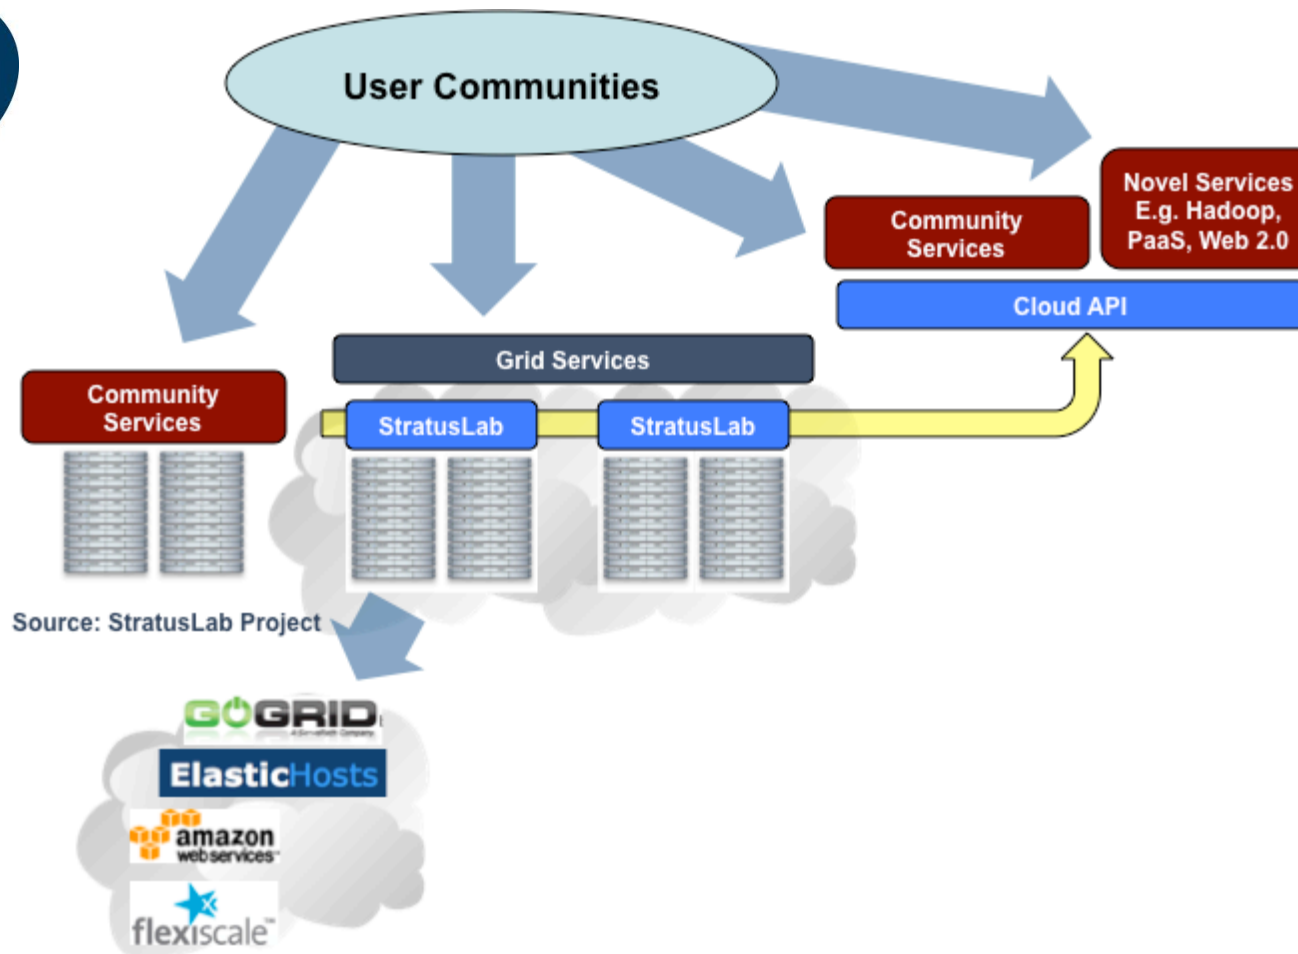
EU grant agreement RI-261552  
**e-Infrastructure**  
(2010-2012)

## 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

---

## Use the Technology and Give us Feedback

---

- Support through several mailing lists
- Report bugs and make feature requests
- Describe your use case in our blog
- Participate in the OpenNebula Technology Days

## 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 workshops or tutorials



## Sponsors

---

### dsa-research.org

- **European Commission:** RESERVOIR (EU agreement 215605), StratusLab (EU agreement 261552), BonFIRE (EU agreement 257386) and 4CaaS (EU agreement 258862).
- **Ministry Science & Innovation:** HPCcloud 2010-2012, MICINN TIN2009-07146
- **Community of Madrid:** MEADIANET 2010-2013 CAM S2009/TIC-1468
- **C12G Labs** dedicates an amount of its own engineering resources to support and develop OpenNebula

## The OpenNebula Community

---

- **The OpenNebula Team:** Ignacio M. Llorente, Ruben S. Montero, Tino Vazquez, Javier Fontan, Jaime Melis, Carlos Martín, Rafael Moreno, Daniel Molina, Borja Sotomayor...
- ... and many **value community contributors** from several organizations

**Your support and contribution are very much appreciated!**