

HEPiX 2009 Spring  
Umeå University  
Thursday 28<sup>th</sup>, May2009

# The OpenNebula Engine for on-Demand Resource Provisioning

Javier Fontán

[dsa-research.org](http://dsa-research.org)

Distributed Systems Architecture Research Group  
Universidad Complutense de Madrid




# What is OpenNebula?

*The OpenNebula Virtual Infrastructure Engine*

## Extending the Benefits of Virtualization to Clusters

- Dynamic deployment and re-placement of virtual machines on a pool of physical resources
- Transform a rigid distributed physical infrastructure into a flexible and agile virtual infrastructure




- Backend of Public Cloud: Internal management of the infrastructure
- Private Cloud: Virtualization of cluster or data-center for internal users
- Cloud Interoperation: On-demand access to public clouds

# Virtual Machine Management Model

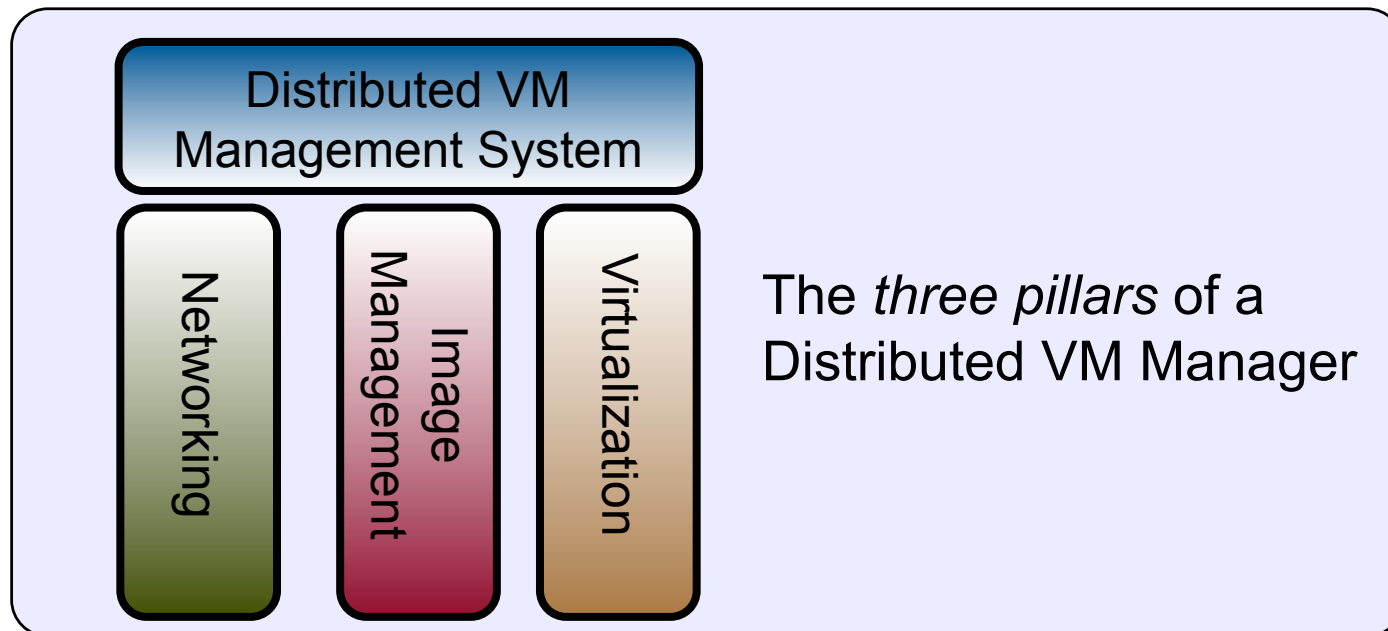
The OpenNebula Virtual Infrastructure Engine

## Service as Management Entity

- Service structure
  - Service components run in VMs
  - Inter-connection relationship
  - Placement constraints
- The VM Manager is service agnostic
- Provide infrastructure context



## Distributed VM Management Model





# Benefits

## *The OpenNebula Virtual Infrastructure Engine*

### System Manager

---

- Centralized management of VM workload and distributed infrastructures
- Support for VM placement policies: balance of workload, server consolidation...
- Dynamic resizing of the infrastructure
- Dynamic partition and isolation of clusters
- Support for heterogeneous workload
- Dynamic scaling of private infrastructure to meet fluctuating demands

### Service Manager

---

- On-demand provision of virtual machines

### System Integrators

---

- Open and flexible architecture and interfaces, open source software
- Integration with any component in the virtualization/cloud ecosystem, such as cloud providers, hypervisors, cloud-like interfaces, virtual image managers, service managers, schedulers...



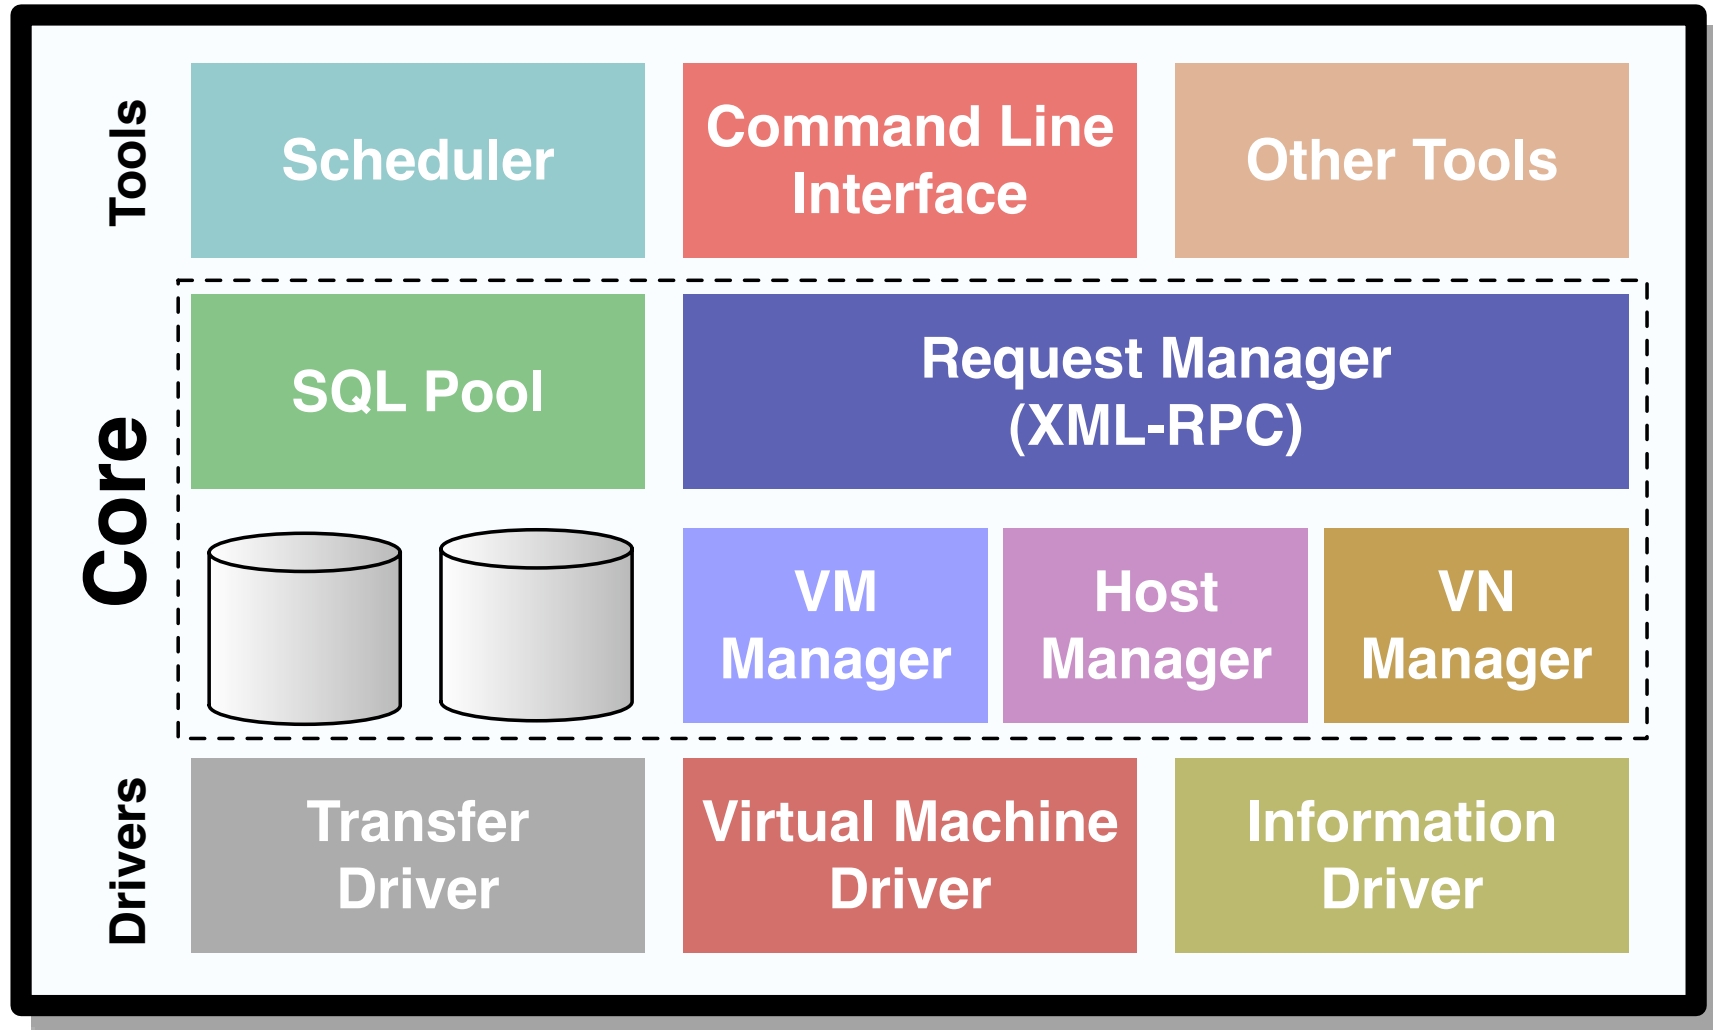
# Features

## *The OpenNebula Virtual Infrastructure Engine*

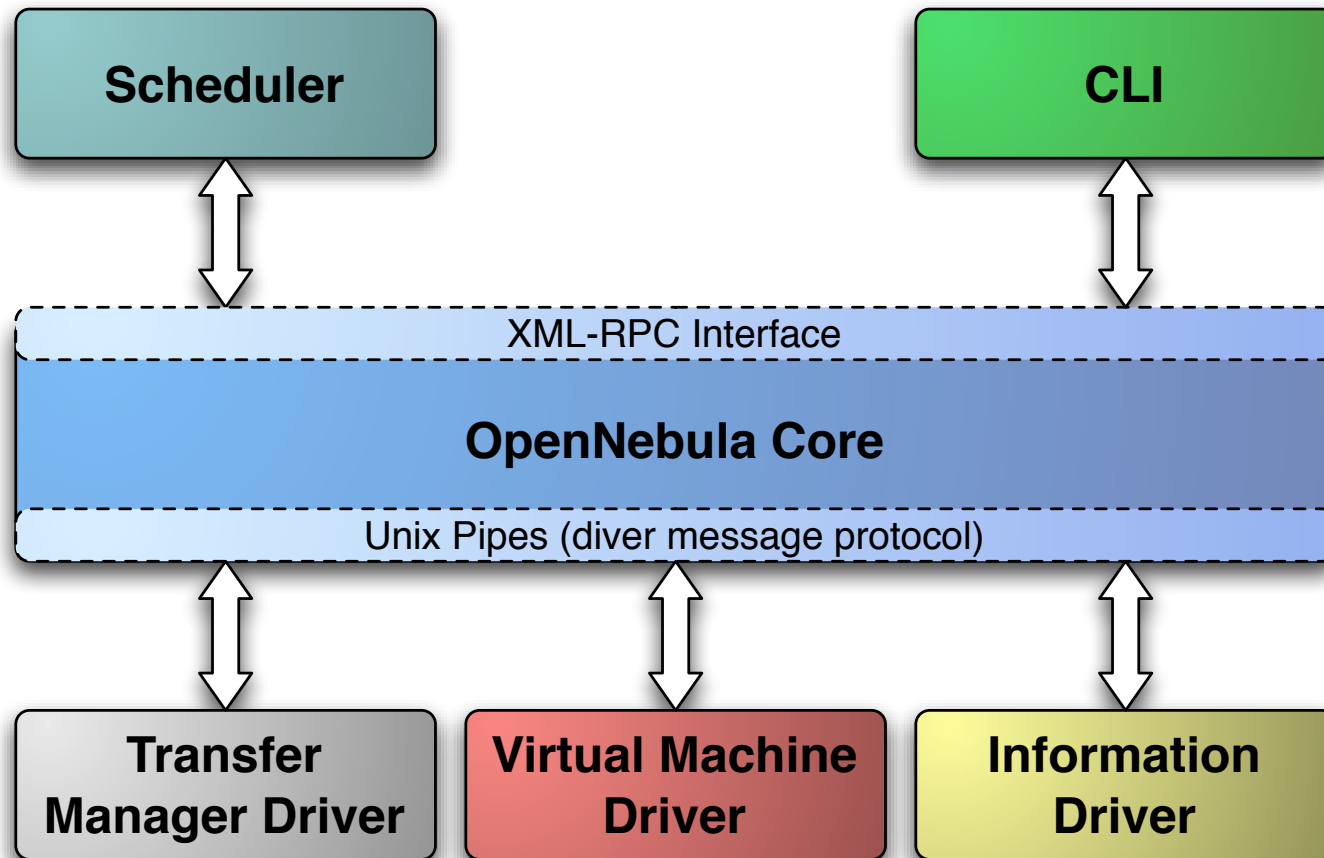
Feature	Function
<b>User Interface</b>	<ul style="list-style-type: none"><li>• Unix-like CLI to manage VM life-cycle and physical boxes</li><li>• XML-RPC API and libvirt interface</li></ul>
<b>Scheduler</b>	<ul style="list-style-type: none"><li>• Requirement/rank matchmaker</li><li>• Generic framework to build any scheduler</li></ul>
<b>Virtualization Management</b>	<ul style="list-style-type: none"><li>• Xen, KVM and libvirt connectors</li><li>• Amazon EC2</li></ul>
<b>Image Management</b>	<ul style="list-style-type: none"><li>• General mechanisms to transfer and clone VM images</li></ul>
<b>Network Management</b>	<ul style="list-style-type: none"><li>• Definition of virtual networks to interconnect VMs</li></ul>
<b>Fault Tolerance</b>	<ul style="list-style-type: none"><li>• Persistent database backend to store host and VM information</li></ul>
<b>Scalability</b>	<ul style="list-style-type: none"><li>• Tested in the management of hundreds of VMs</li></ul>
<b>Installation</b>	<ul style="list-style-type: none"><li>• Installation on a UNIX cluster front-end without requiring new services in the remote resources</li><li>• Distributed in Ubuntu 9.04 (Jaunty Jackalope)</li></ul>

# Open and Flexible Architecture

*The OpenNebula Virtual Infrastructure Engine*



# Process Separation



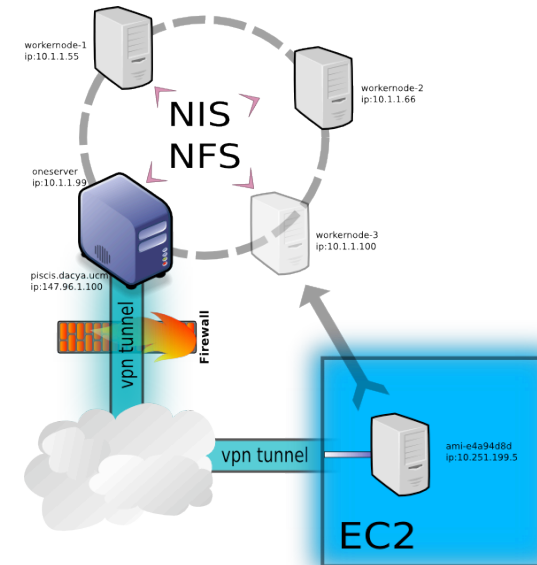
- Scheduler is a separated process, just like command line interface.
- Drivers are also separated processes using a simple text messaging protocol to communicate with OpenNebula Core Daemon (oned)

# Use Cases

## The OpenNebula Virtual Infrastructure Engine


### On-demand Scaling of Computing Clusters

- Elastic execution of a SGE computing cluster
- Dynamic growth of the number of worker nodes to meet demands using EC2
- Private network with NIS and NFS
- EC2 worker nodes connect via VPN



### On-demand Scaling of Web Servers

- Elastic execution of the NGinx web server
- The capacity of the elastic web application can be dynamically increased or decreased by adding or removing NGinx instances





# Ecosystem

## The OpenNebula Virtual Infrastructure Engine

### Schedulers

- Haizea: Open-source VM-based lease management architecture (allows AR of capacity).

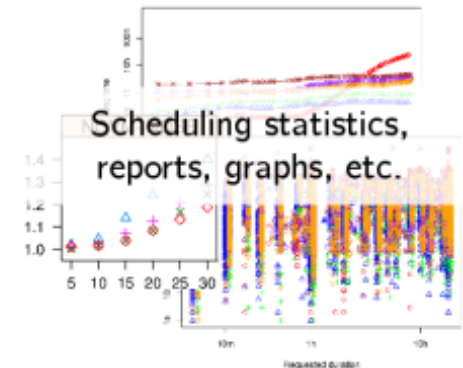



Lease requests  
 "I need 10 nodes, each with  
 2 CPUs, 4GB of memory,  
 from 2pm to 4pm"



With  
OpenNebula

Simulation



### Interfaces

- Libvirt: Provides an abstraction of a whole cluster of resources as one host, hiding specific hypervisor details.
- Nimbus: Can be used as a WSRF or EC2 front-end.

### Plug-Ins

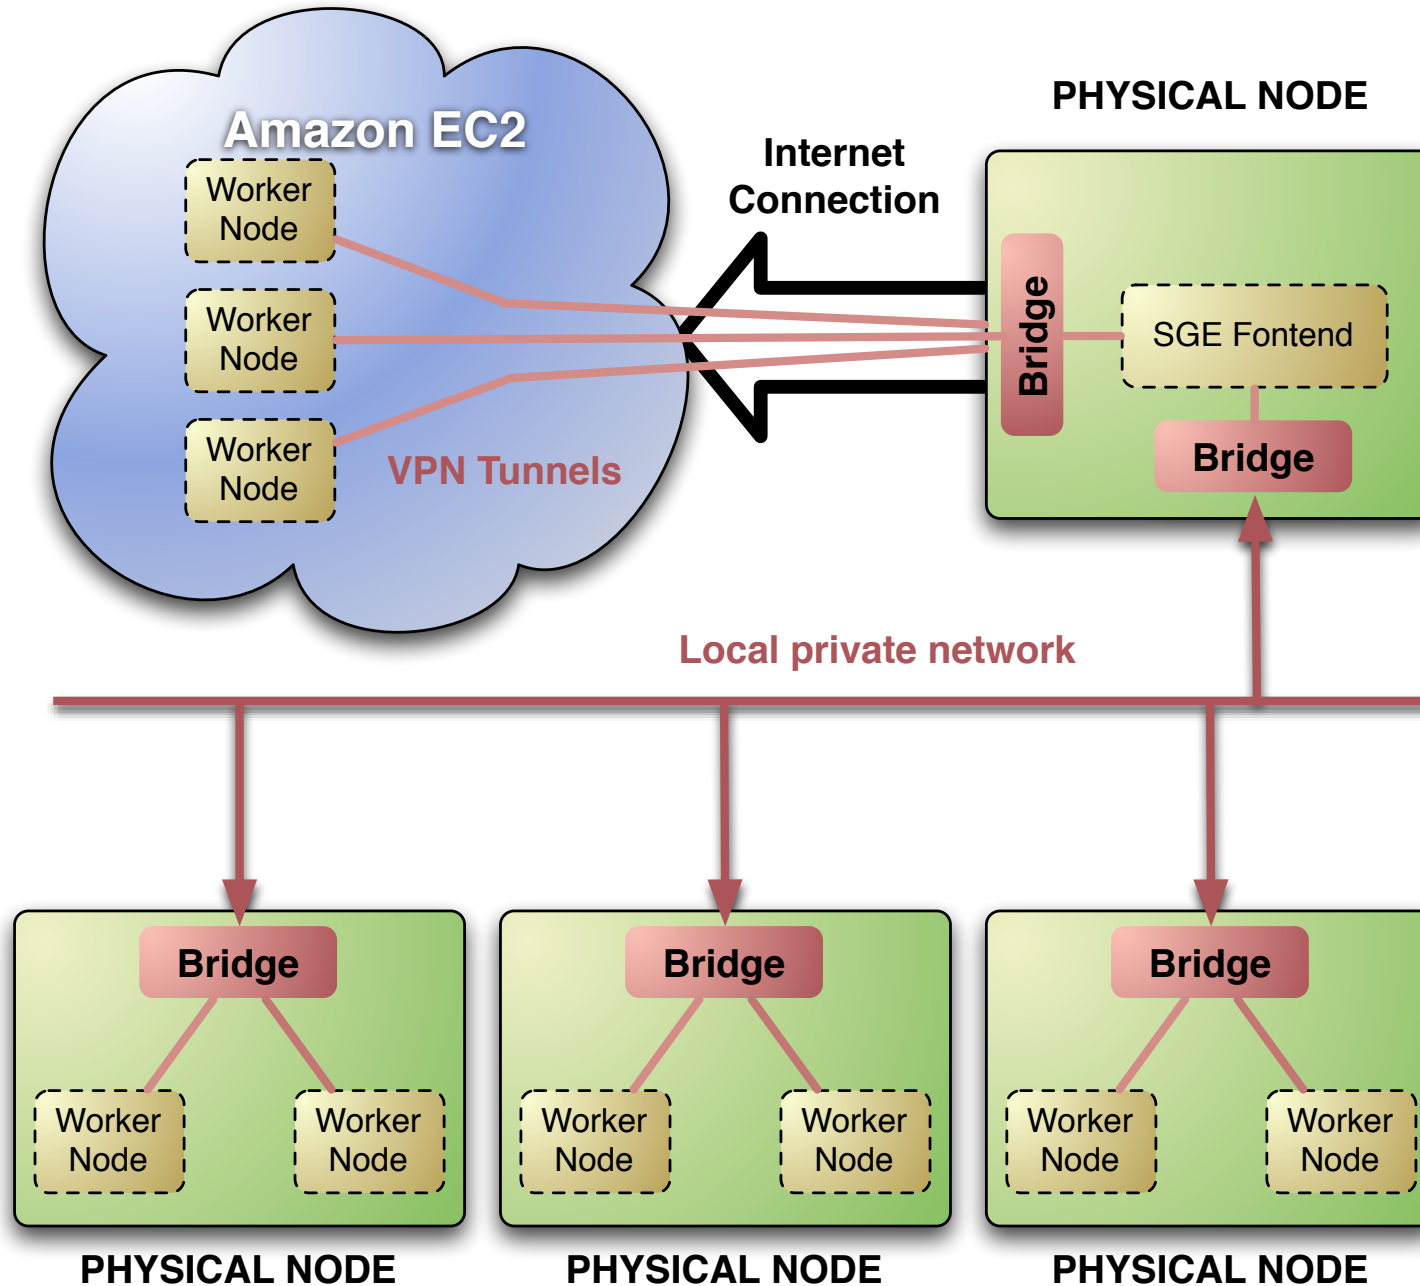
- ElasticHosts: Enables the dynamically increase capacity of your virtualized infrastructure to meet fluctuating peak demands using a cloud provider.



# Main Differences With Other Open Source VM Managers

- Mainly destined to SysAdmins as the tools provided to manage VM's and Physical Hosts were developed by and for SysAdmins. Command Line interface is similar to LRMS and an API exists for easy scriptability.
- Easy customization of every aspect of VM management. Host monitoring, image transferring and VM management can be easily modified or extended with simple scripts. No modification to the core is needed.
- Multiple hypervisor technologies support. Stable version comes with Xen and KVM drivers. New drivers can be easily created with the provided framework or from scratch. No compilation needed. Next version will come with VMWare support.
- Externalizing peak loads. OpenNebula is capable of managing VMs that are not located in local resources. To do this EC2 and ElasticHosts drivers are provided.
- Contextualization and advanced configuration. OpenNebula is able to generate iso images with that contains any file you may need for the VM. Hooks can be provided so you can start your own actions when a VM state is changed.

# SGE Deployment Architecture



# The OpenNebula VM Manager

THANK YOU FOR YOUR ATTENTION!!!  
More info, downloads, mailing lists at  
[www.OpenNebula.org](http://www.OpenNebula.org)

OpenNebula is partially funded by the “RESERVOIR– Resources and Services Virtualization without Barriers” project  
EU grant agreement 215605



Acknowledgements [www.reservoir-fp7.eu/](http://www.reservoir-fp7.eu/)

- Ignacio M. Llorente
- Rubén S. Montero
- Raúl Sampedro
- Tino Vazquez
- Rafael Moreno