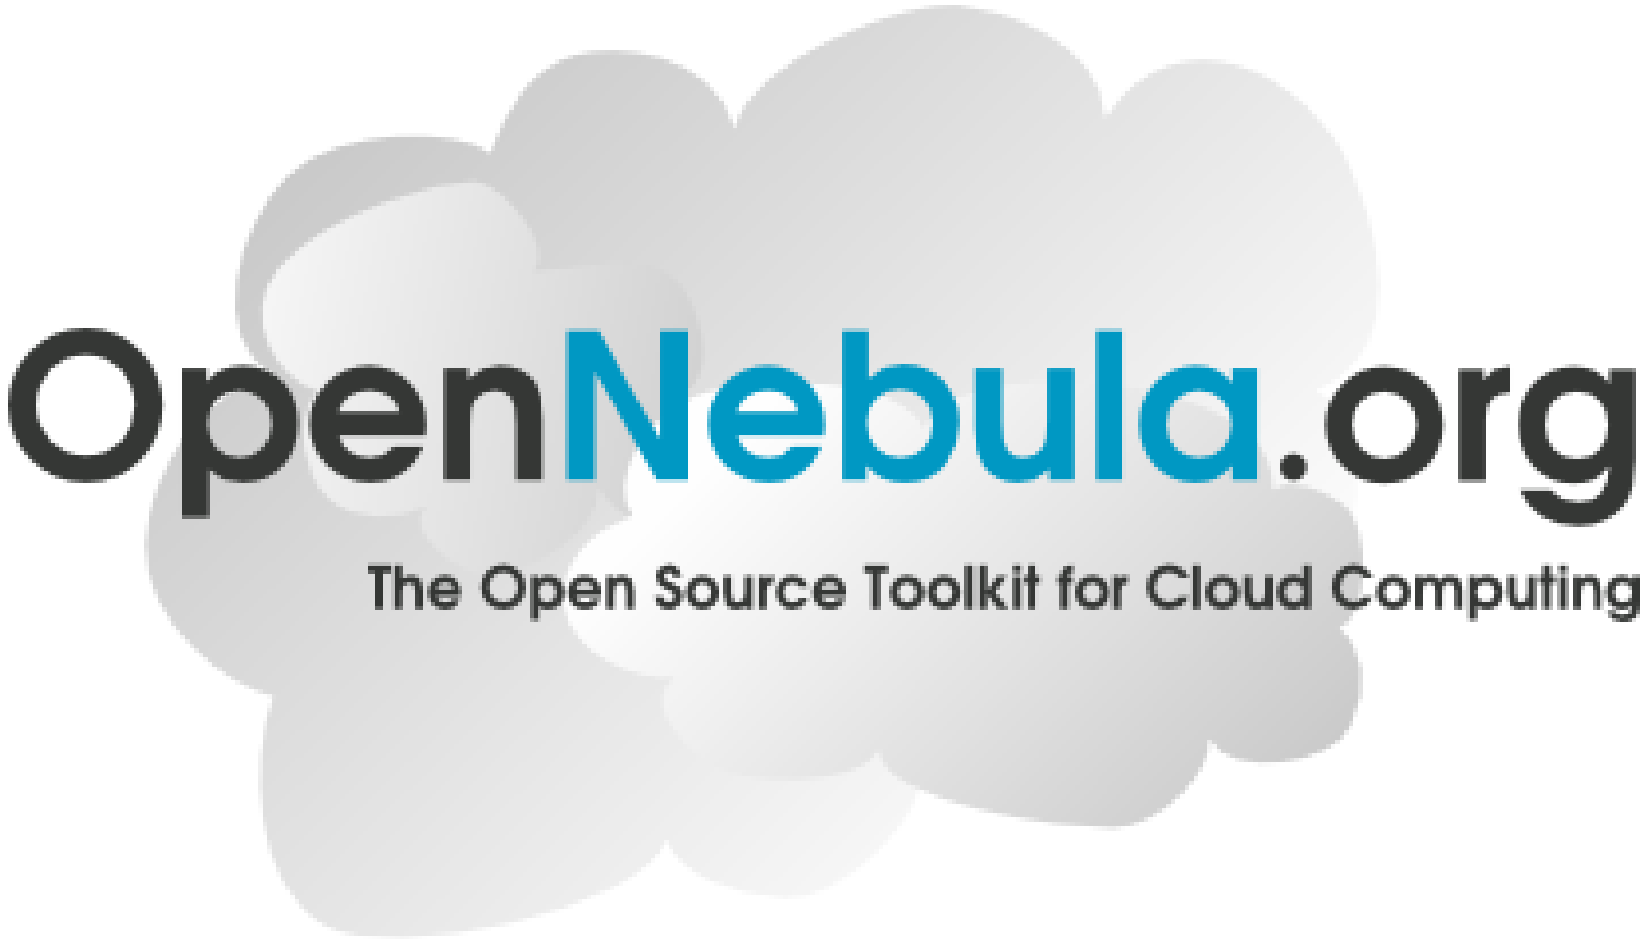
OpenNebula.org
The Open Source Toolkit for Cloud Computing

The OpenNebula Cloud Toolkit: Experiences and Outlook

Borja Sotomayor
University of Chicago

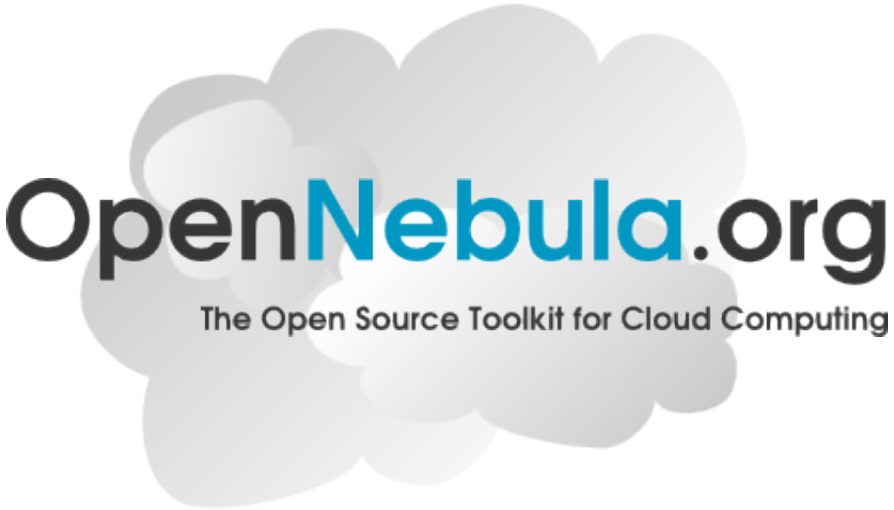

What is OpenNebula?

Experiences

Ecosystem

Outlook

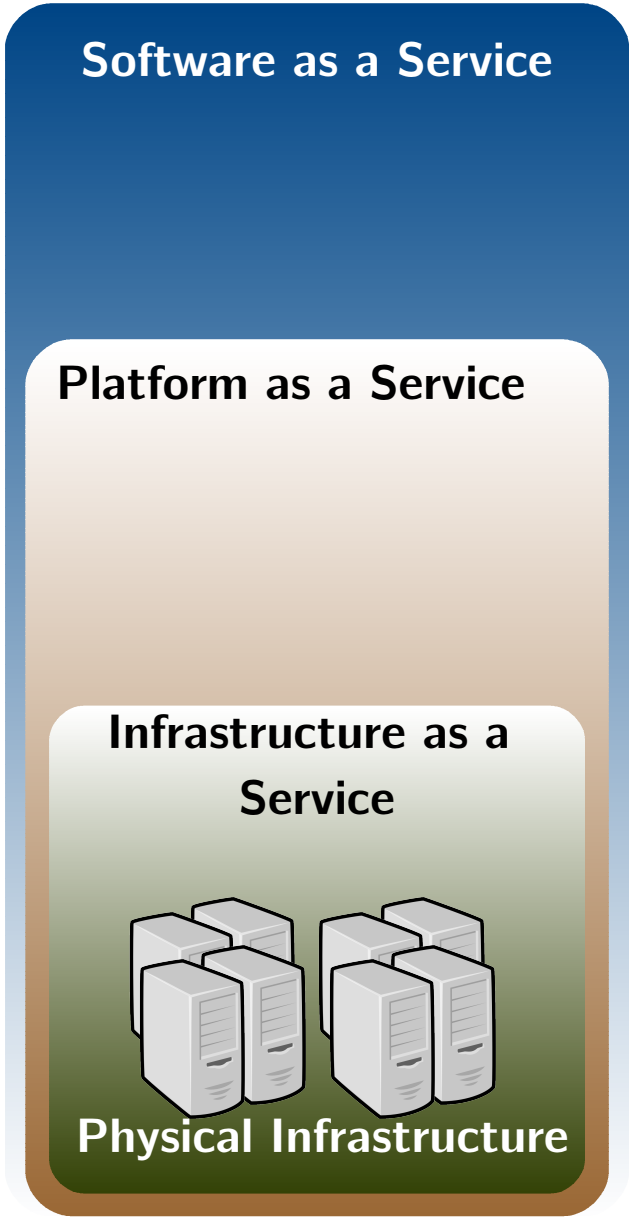

skype

facebook

Google docs

Windows Azure

force.com platform as a service

GOGRID

flexiscale

amazon web services

$0.085 / machine / hour

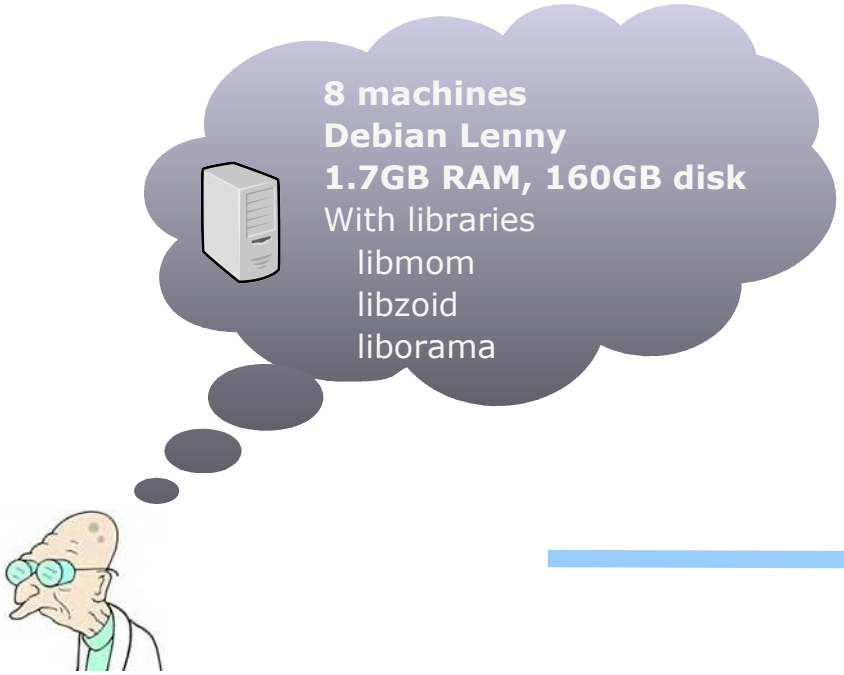

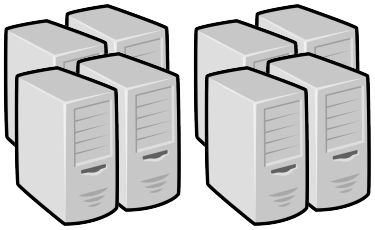

amazon web services™

http://aws.amazon.com/ec2/

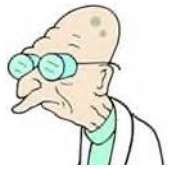

IaaS

Public Interface

Virtual Infrastructure Manager

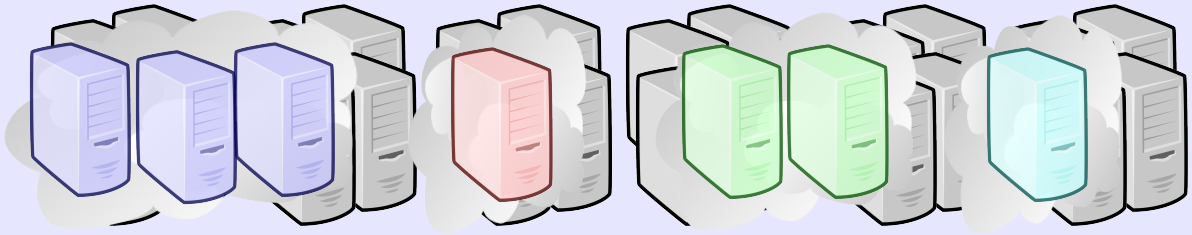

Physical resources with Virtual Machine Managers (Xen, KVM, VMWare, etc.)

IaaS

Public Interface

Virtual Infrastructure Manager

Resource scheduling

Contextualization

Storage management

Virtual network management

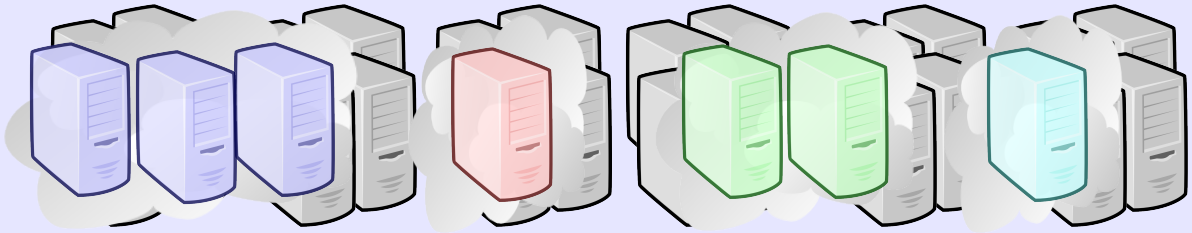

Physical resources with Virtual Machine Managers (Xen, KVM, VMWare, etc.)

Cool!

. . . if you're Amazon

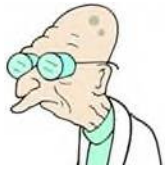

External Users

Internal Users

Public Interface

Internal Interface

Private Cloud

Hybrid Cloud

Virtual Infrastructure Manager

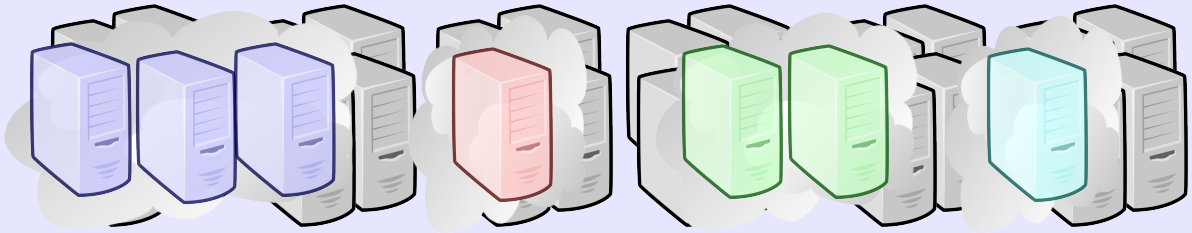

Physical resources with Virtual Machine Managers (Xen, KVM, VMWare, etc.)

External Cloud

The OpenNebula Cloud Toolkit: Experiences and Outlook

Borja Sotomayor
University of Chicago

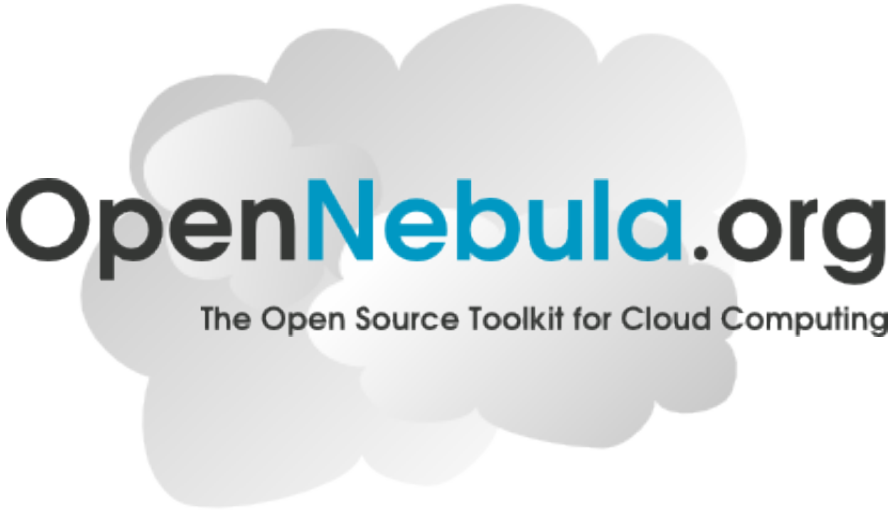

What is OpenNebula?

Experiences

Ecosystem

Outlook

~~OpenNebula is a virtual infrastructure manager~~ *

OpenNebula is a standards-based
open-source **toolkit** to build private,
public and hybrid clouds.

Most development takes place at the University Complutense of Madrid, and is funded by several European and Spanish grants.
Project has been ongoing since 2005.

* Managing VMs is a big part of building an IaaS cloud, but not the only part.

Why a “toolkit”?

End-User

Wants cloud interfaces to manage virtual machines, network and storage.

Prefers popular interfaces like Amazon EC2

Multi-tier services as a basic management entity

Wants cloudbursting to public clouds and possibly to partner clouds.

Cloudbursting must be transparent to users

Sysadmin

Wants administration interface.

Needs control over resource allocation policies

May have to support existing data center services.

Needs to integrate with products and services in the virtualization/cloud ecosystem such as cloud providers, hypervisors, virtual image managers, service managers, management tools, schedulers...

Must be easy to add new functionality and to embed into other platforms.

One solution does not fit all requirements and constraints.

The OpenNebula design philosophy

There cannot be turnkey solutions for IaaS clouds,
so OpenNebula shouldn't aim to be one.

First and foremost, provide an architecture that is open,
flexible, and extensible that allows multiple components to
be orchestrated.

Provide some components of our own, but allow them to
be easily replaceable by others.

Seriously, though, what *is* OpenNebula?

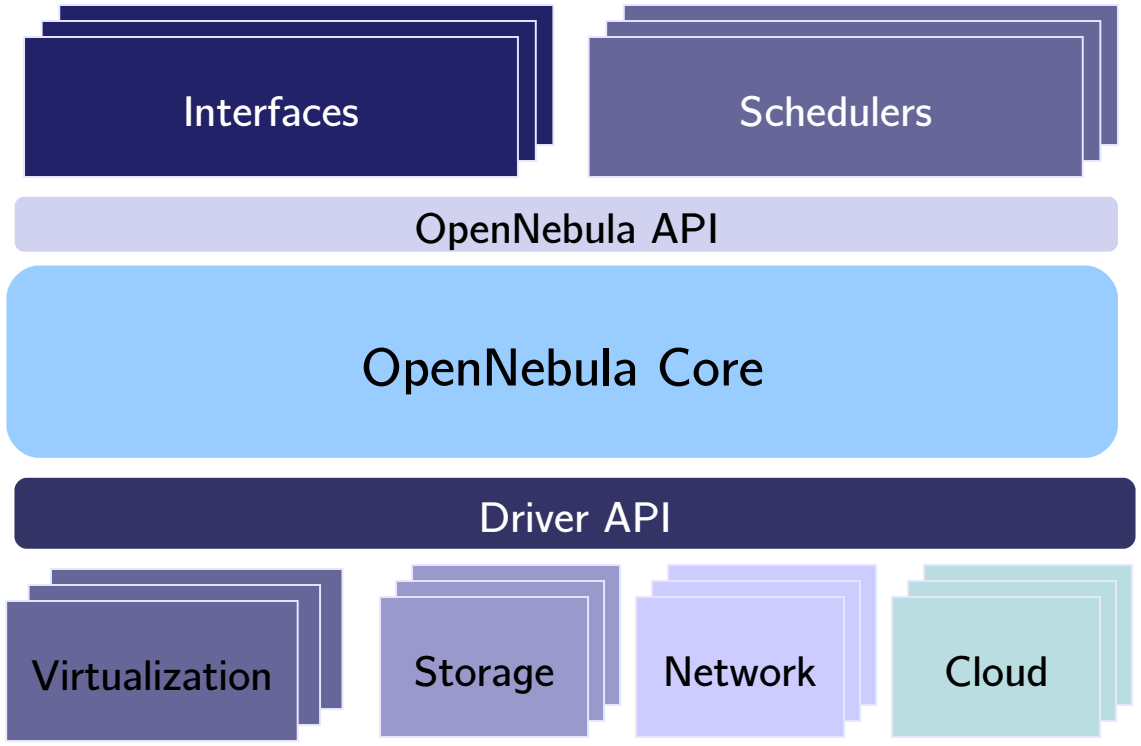
Interfaces
Schedulers
OpenNebula API
OpenNebula Core
Driver API
Virtualization
Storage
Network
Cloud

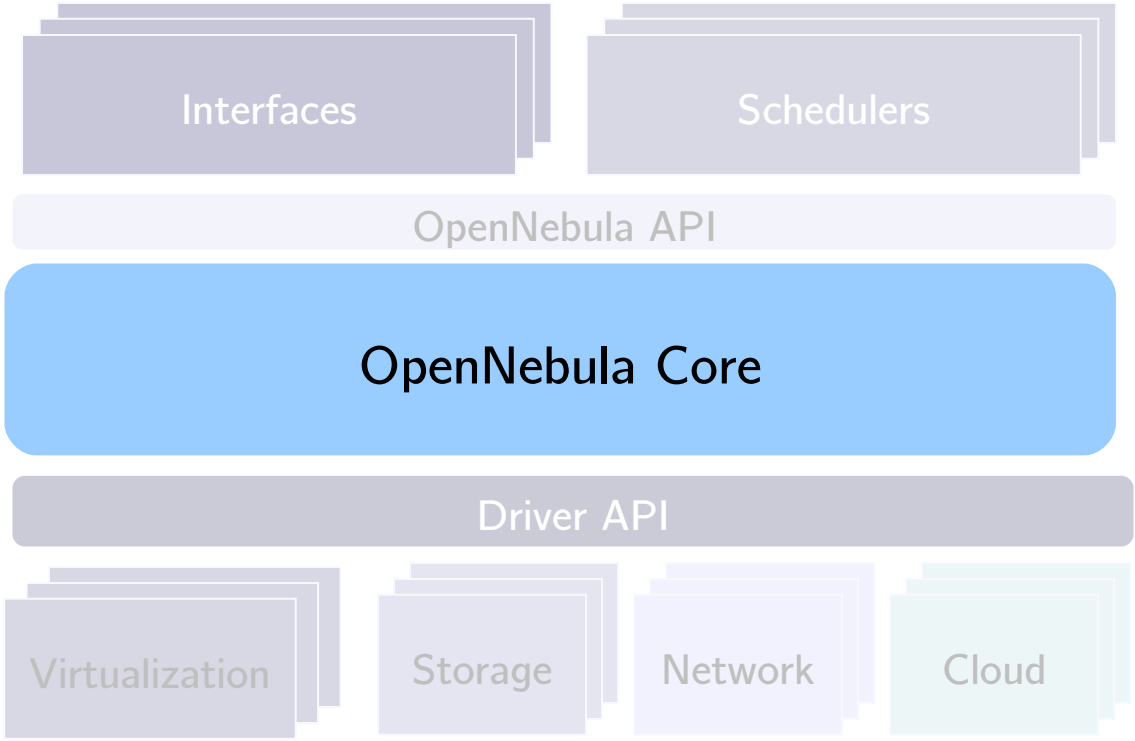

Handles orchestration of all the different components. Also handles some cross-cutting features, such as user management, persistence, etc.

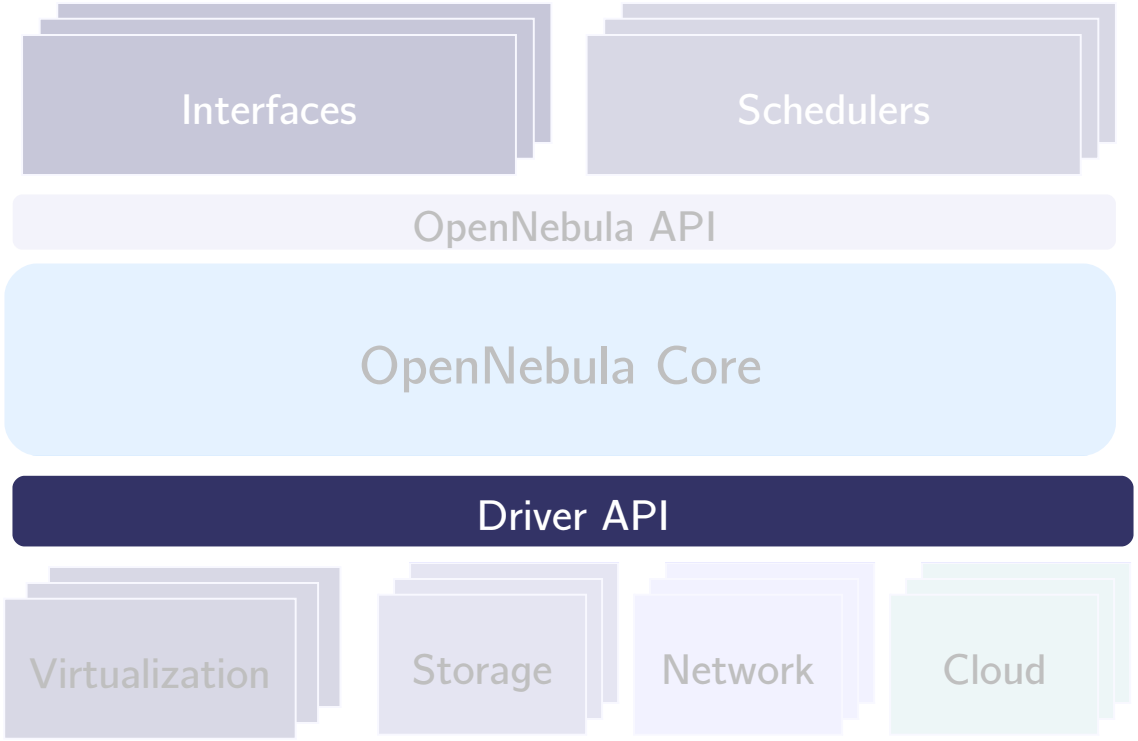

Provides a layer of abstraction over lower-level operations.

Drivers are self-contained and can be written without modifying OpenNebula's core.

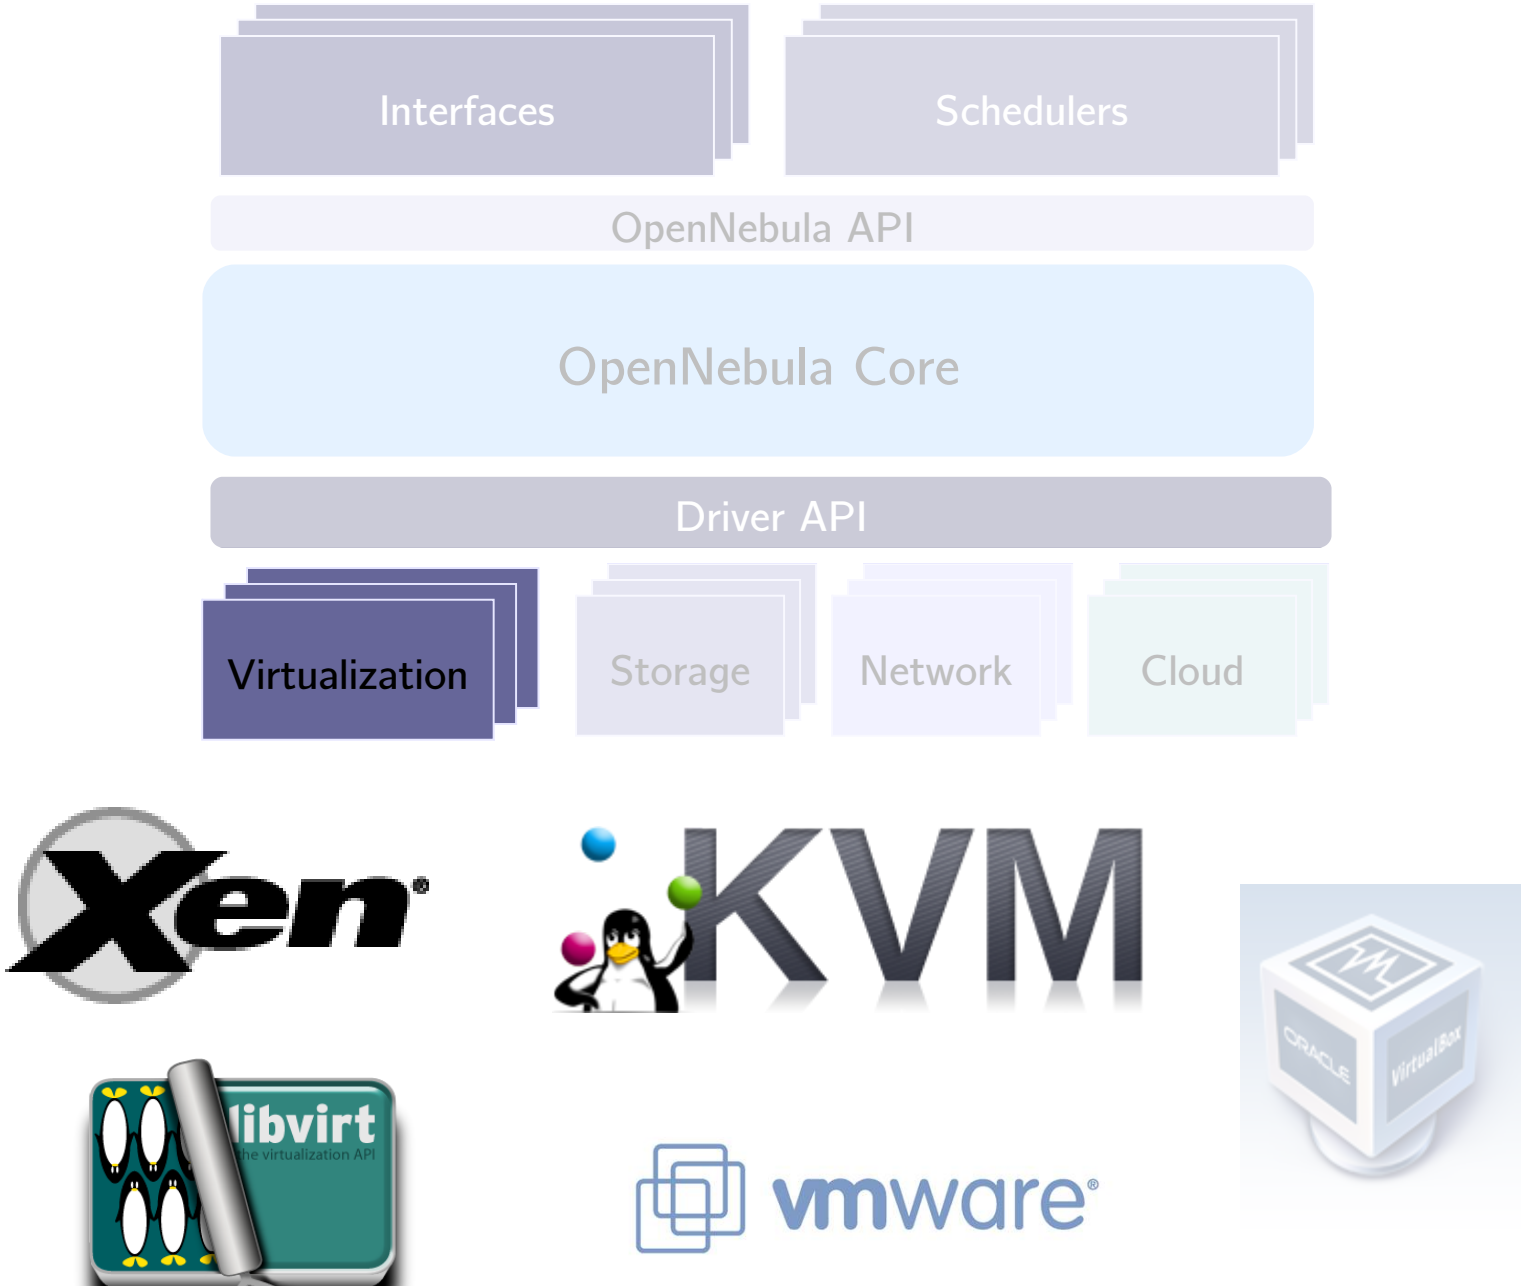
Interfaces
Schedulers
OpenNebula API
OpenNebula Core
Driver API
Virtualization
Storage
Network
Cloud
Xen
KVM
libvirt
the virtualization API
vmware
ORACLE
VirtualBox

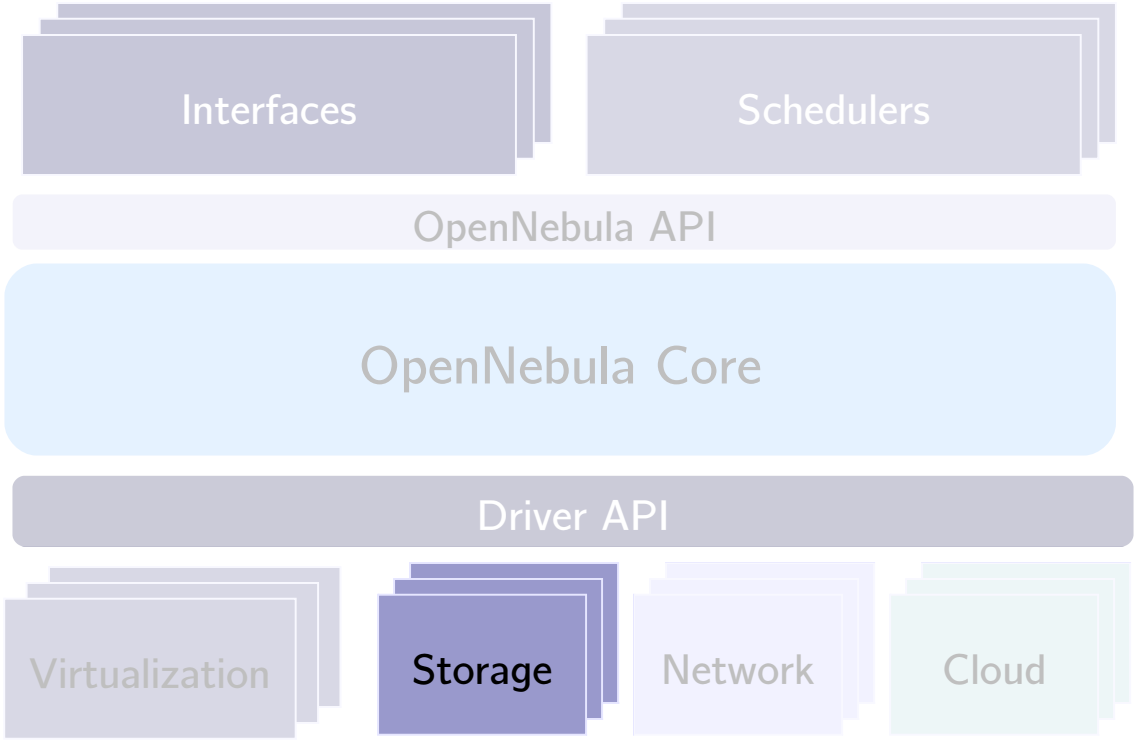

Images on shared NFS

SCP from image repository

LVM

Contextualization of disk images

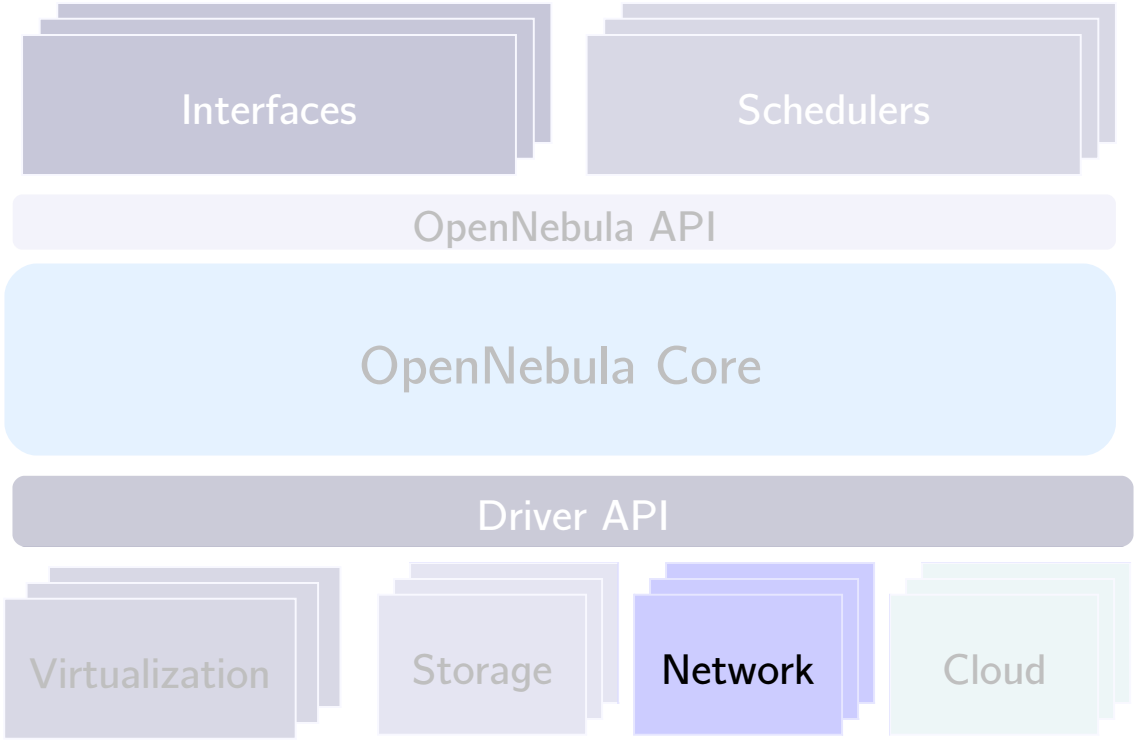

Creation of virtual networks

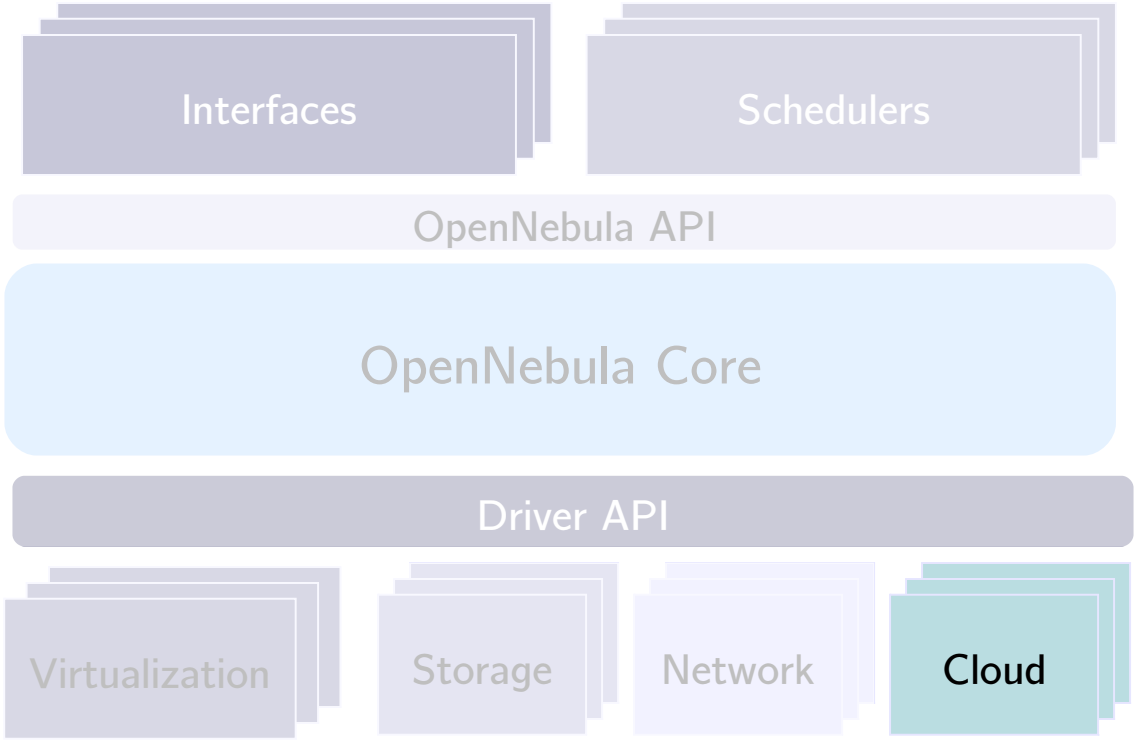

amazon web services™

ElasticHosts
Flexible servers in the cloud

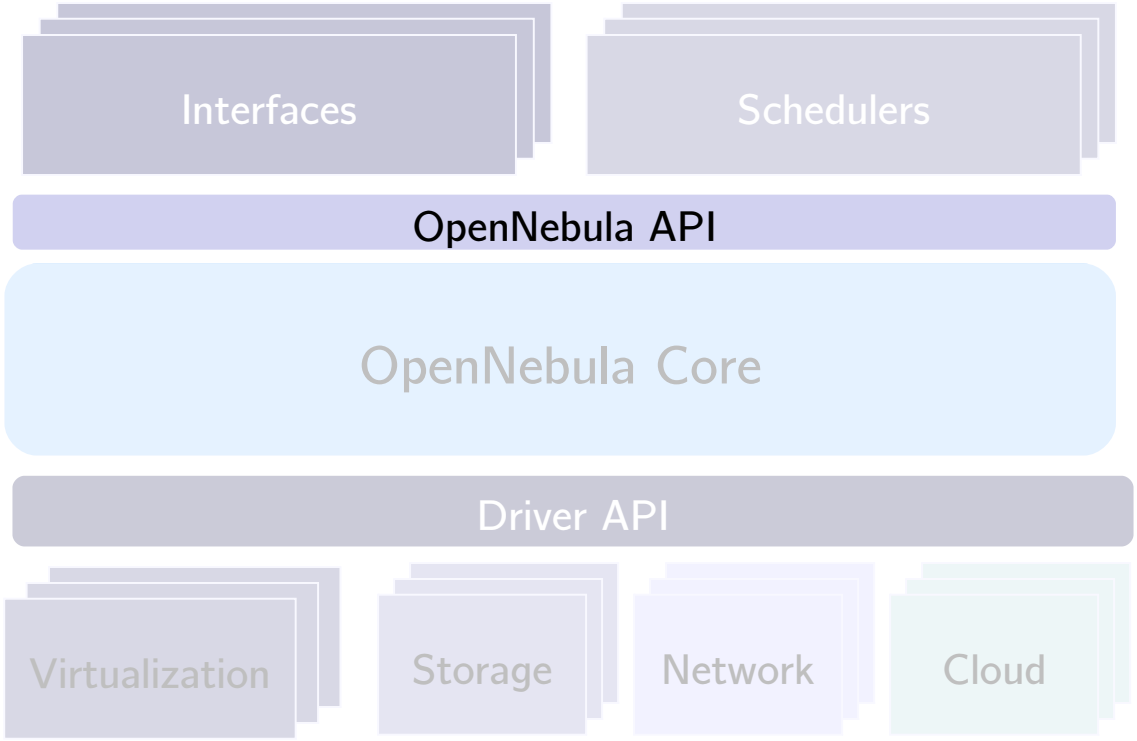

XML-RPC CLI OpenNebula Cloud API

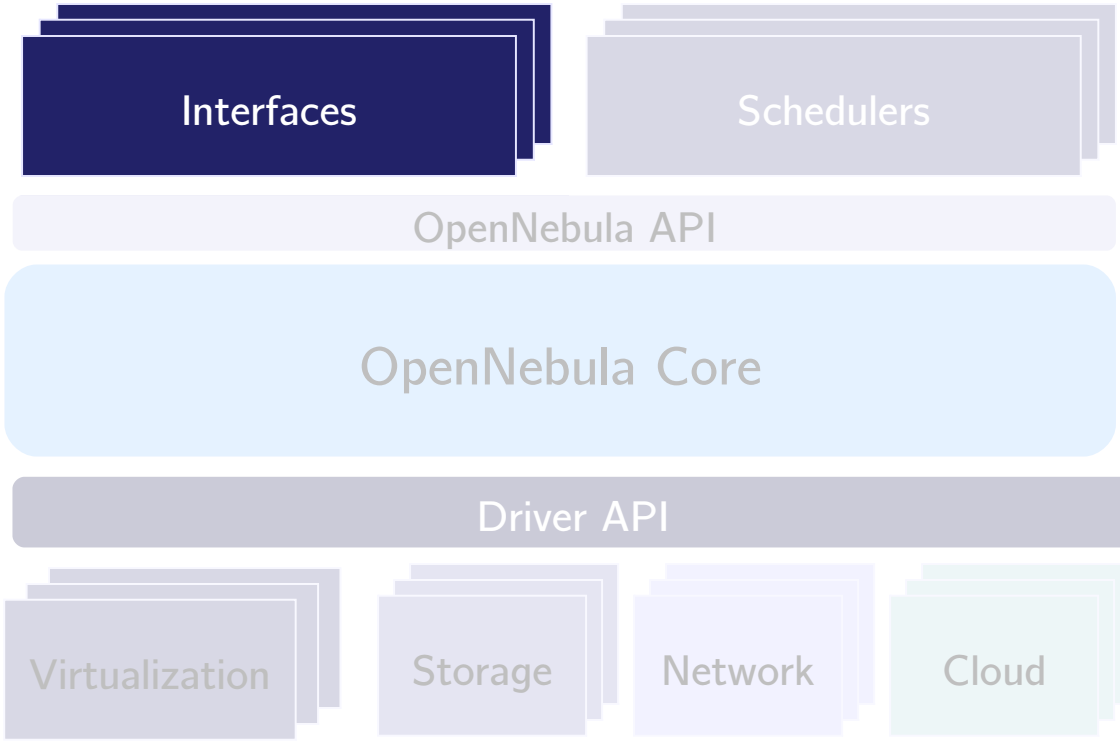

OCCI
OPEN CLOUD COMPUTING INTERFACE

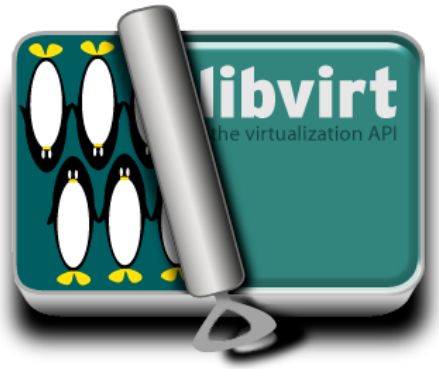

amazon
web services™

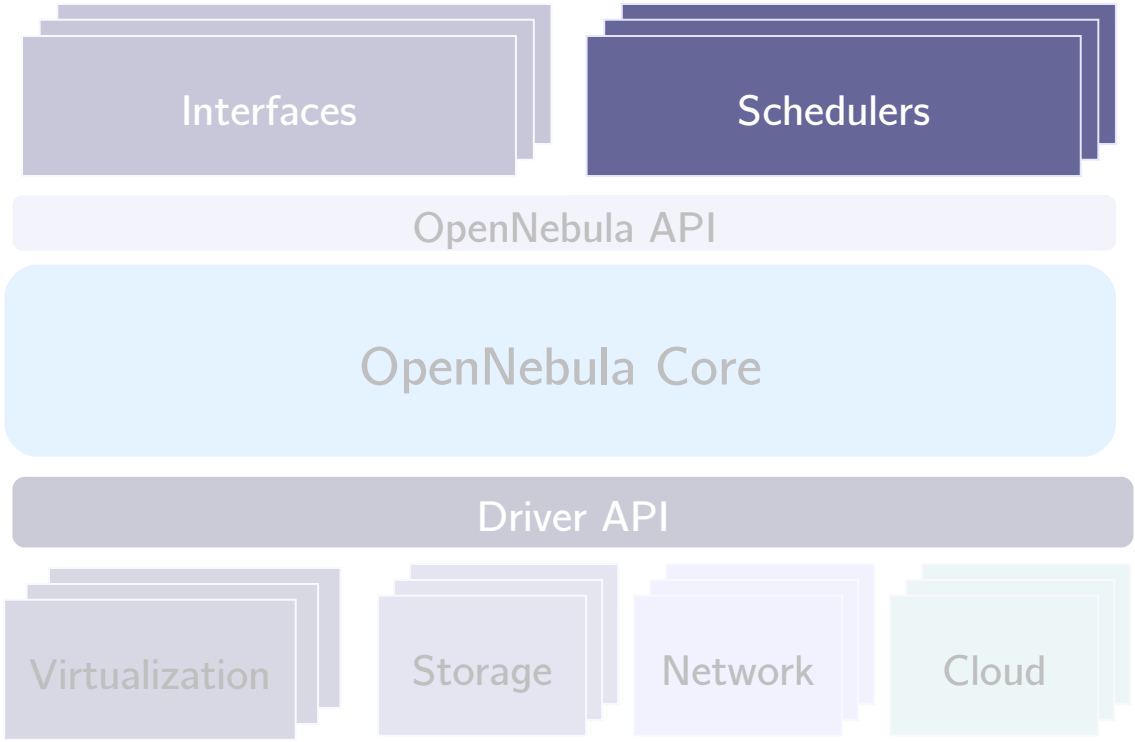

Matchmaking scheduler with a configurable ranking policy

Geared towards immediate scheduling, with basic queueing

How does it compare to other solutions?

BiG Grid, *Virtualization of worker nodes,* Working group progress report. 02/02/2010 http://tinyurl.com/big-opennebula

MIDDLEWARE 2009

Cerbelaud, D., Garg, S., and Huylebroeck, J. *Opening the clouds: qualitative overview of the state-of-the-art open source VM-based cloud management platforms*. In Proceedings of the 10th ACM/IFIP/USENIX international Conference on Middleware 2009

B.Sotomayor, R.Santiago Montero, I.Martín Llorente, I.Foster, *Virtual Infrastructure Management in Private and Hybrid Clouds*. IEEE Internet Computing, vol. 13, no. 5, pp. 14-22, Sep./Oct. 2009.

The OpenNebula Cloud Toolkit: Experiences and Outlook

Borja Sotomayor
University of Chicago

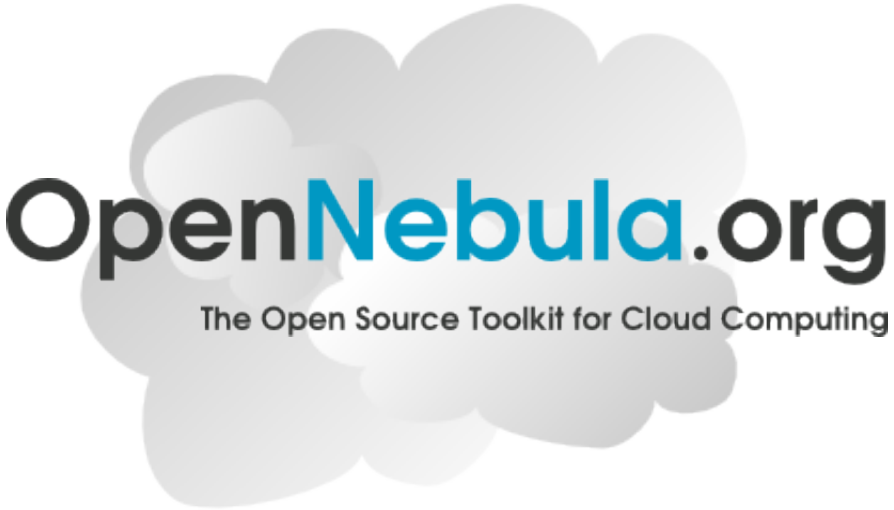

A team at Clemson University and CERN has used OpenNebula to deploy up to 7,500 VMs on 400 physical hosts running Xen.

These VMs are used to run batch jobs (submitted via WLCG CE and managed by LSF)

Used XML-RPC API to add certain autonomic functionality that was important to them, and to integrate with CERN's Quattor (http://www.quattor.org/)

Created, and contributed, drivers for using LVM-based disk images.

More details at http://tinyurl.com/cern-opennebula

The D-Grid (national German Grid initiative) Resource Center Ruhr (DGRZR) has used OpenNebula to manage 247 Blades with a total of 1,972 cores.

Entire D-Grid software stack is run on VMs. Grid worker nodes currently managed with OpenNebula, frontend nodes to follow shortly.

More details at http://tinyurl.com/dgrid-opennebula

SARA is the Dutch National High Performance Computing and e-Science Support Center, and the Dutch supernode in the international Science Grid.

They are currently developing an HPC cloud that uses OpenNebula. Users get their own 'Virtual Private HPC Cluster'

Starting with 128 cores across 16 physical machines running KVM.

Users use a management console developed at SARA to request a new VM. Templates are provided, but users can also configure their own

More details at http://tinyurl.com/sara-opennebula

The BiG Grid Virtual Machine Working Group (in NIKHEF) did an evaluation of several cloud solutions, and recommended using OpenNebula for managing worker node VMs in BiG Grid.

More details at http://tinyurl.com/big-opennebula

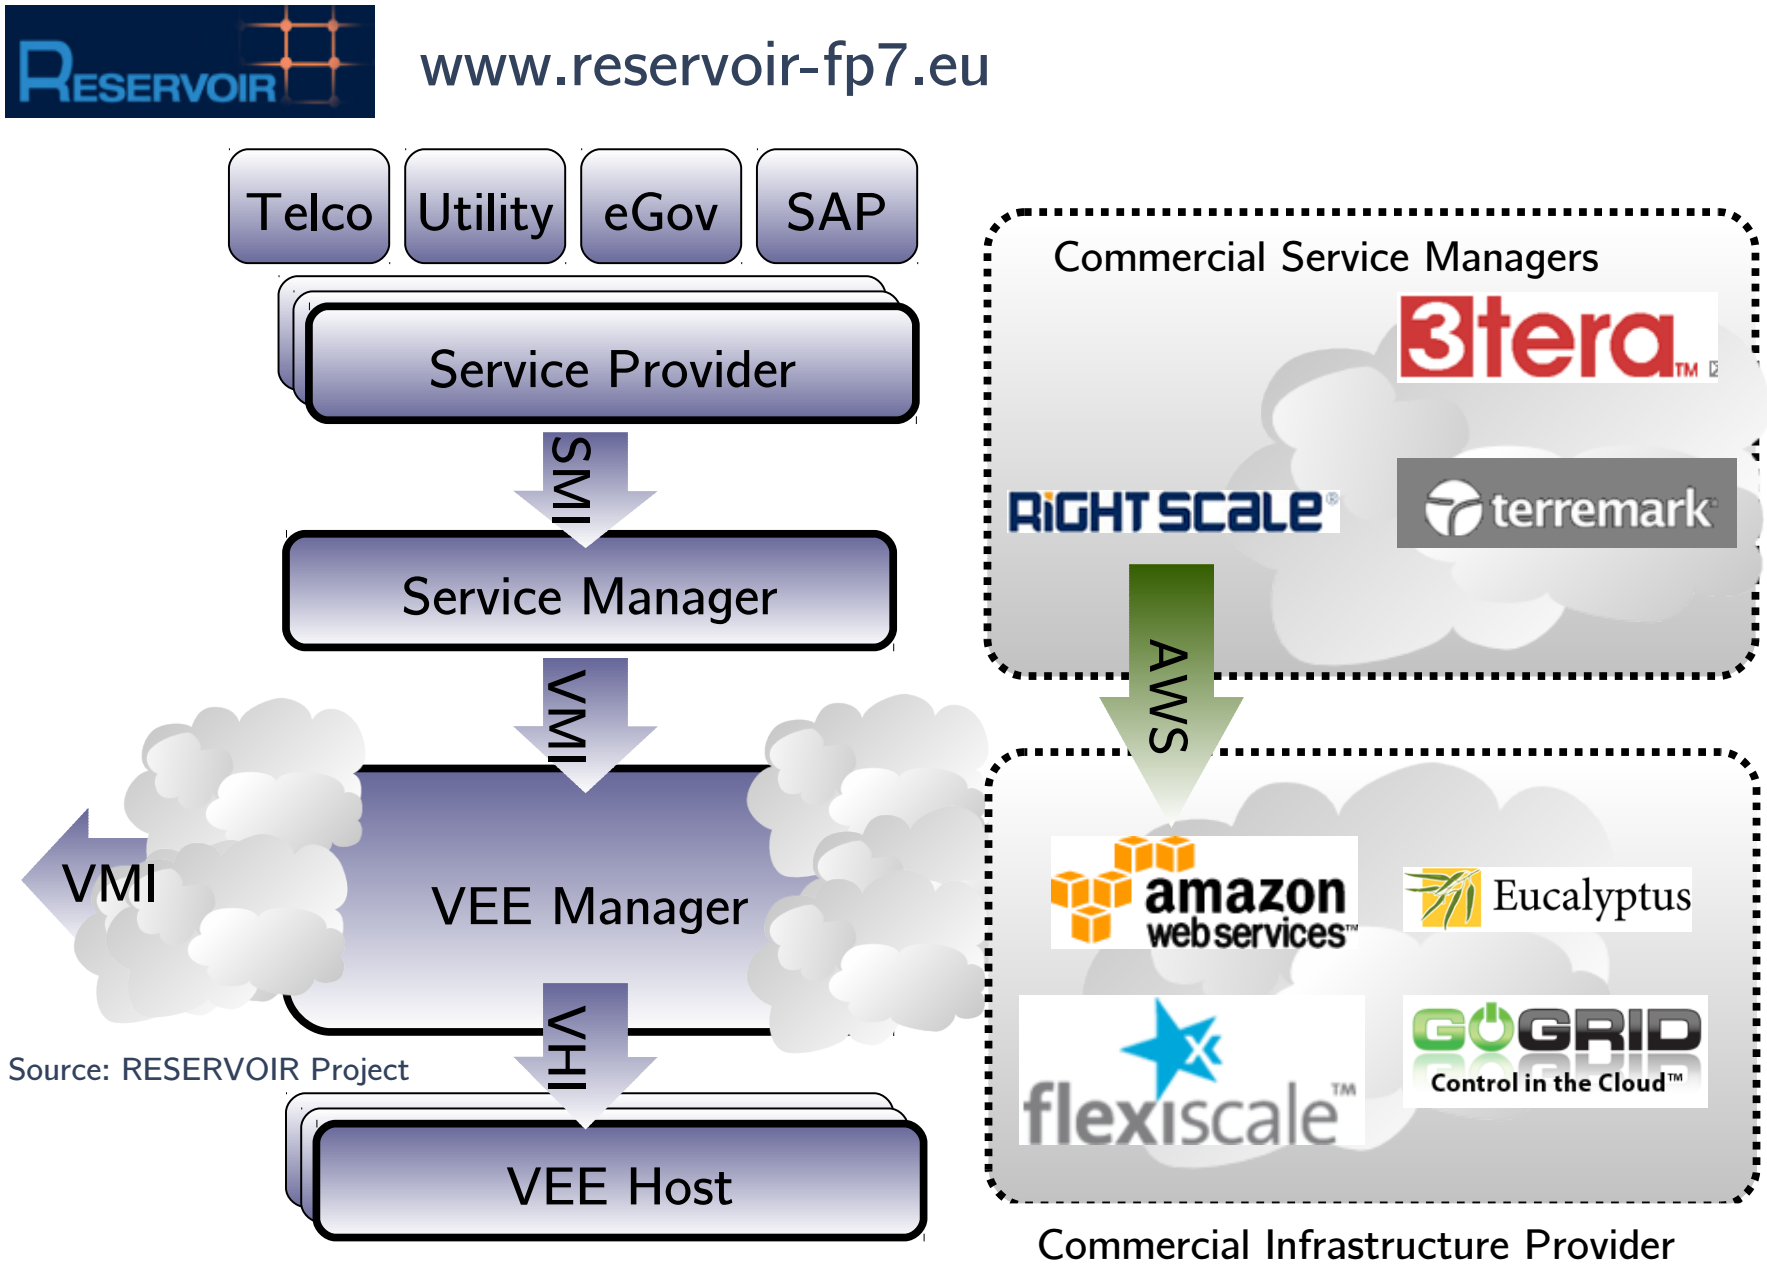
RESERVOIR
www.reservoir-fp7.eu
Telco
Utility
eGov
SAP
Service Provider
SMI
Service Manager
VMI
VMI
VEE Manager
VHI
VEE Host
Source: RESERVOIR Project
Commercial Service Managers
3tera
RIGHTSCALE
terremark
AWS
amazon web services
Eucalyptus
flexiscale
GOGRID
Control in the Cloud
Commercial Infrastructure Provider

Research with OpenNebula

Cloud architectures, federation, interoperability

- 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)
- 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)

High Performance Computing (HPC) clouds

- R. Moreno, R. S. Montero, e I. M. Llorente, *Elastic Management of Cluster-based Services in the Cloud,* First Workshop on Automated Control for Datacenters and Clouds (ACDC09)
- I. M. Llorente, R. Moreno-Vozmediano, and R. S. Montero, *Cloud Computing for On-Demand Grid Resource Provisioning*, Advances in Parallel Computing, Volume 18 (2009): "High Speed and Large Scale Scientific Computing", pp. 177 - 191. IOS Press, 2009.
- B. Sotomayor, R. S. Montero, I. M. Llorente and I. Foster, *Resource Leasing and the Art of Suspending Virtual Machines*, IEEE International Conference on High Performance Computing and Communications (HPCC-09), Seoul, Korea

Research with OpenNebula

Service Management

- Luis Rodero-Merino, Luis M. Vaquero, Victor Gil, Fermín Galán, Javier Fontán, Rubén S. Montero, and Ignacio M. Llorente, *From infrastructure delivery to service management in clouds*, Future Generation Computer Systems. In press

Energy-efficient cloud computing

- G. von Laszewski, L. Wang, A. J. Younge, X. He, *Power-Aware Scheduling of Virtual Machines in DVFS-enabled Clusters*, Proceedings of IEEE International Conference on Cluster Computing and Workshops, 2009. CLUSTER '09.

OpenNebula in industry

CloudScaling (http://cloudscaling.com/)
According to Randy Bias, CEO, "*Cloudscaling has had great success with OpenNebula. Unlike many of the other open source virtual infrastructure management tools, ONE is cleanly written, modular, and easily extensible. We use it regularly in our labs and in some client engagements. Highly recommended.*"

Morph Labs (http://www.mor.ph/) uses OpenNebula in its mCloud Controller product.

The OpenNebula Cloud Toolkit: Experiences and Outlook

Borja Sotomayor
University of Chicago

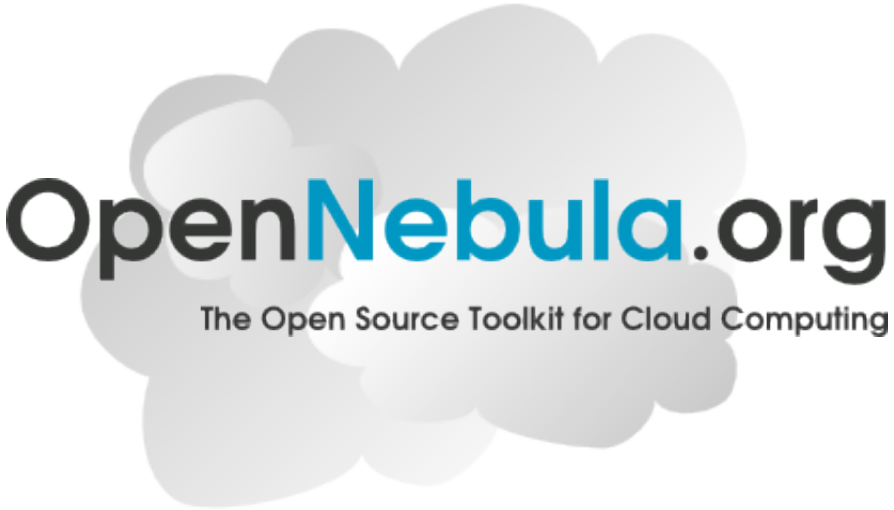

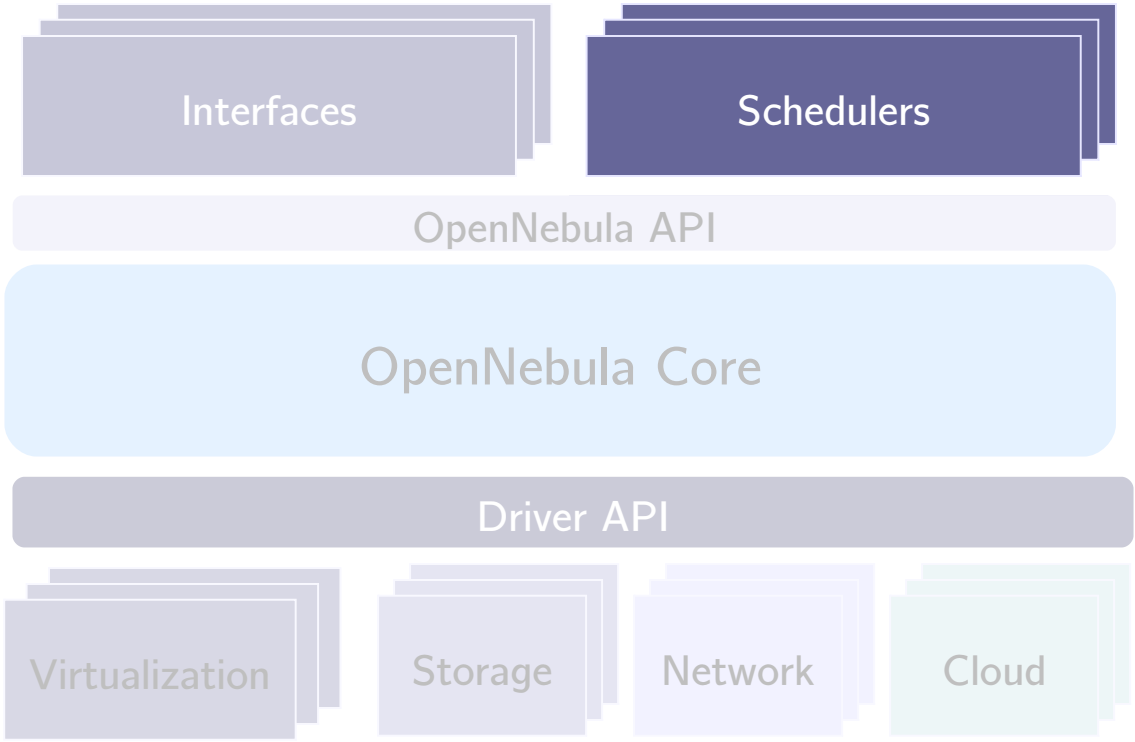

HAIZEA

More featureful scheduler, including support for advance reservation, queueing, and pluggable scheduling policies.
Developed at the University of Chicago
http://haizea.cs.uchicago.edu/

Haizea is a lease manager that can act as a scheduling backend for OpenNebula, providing advanced functionality such as:

- Advance reservation of capacity
- Best-effort scheduling with backfilling
- Resource preemption (using VM suspend/resume/migrate)
- Policy engine, allowing developers to write pluggable scheduling policies in Python
- Includes a simulation mode (useful for researchers testing scheduling algorithms)

Haizea research publications: http://haizea.cs.uchicago.edu/pubs.html

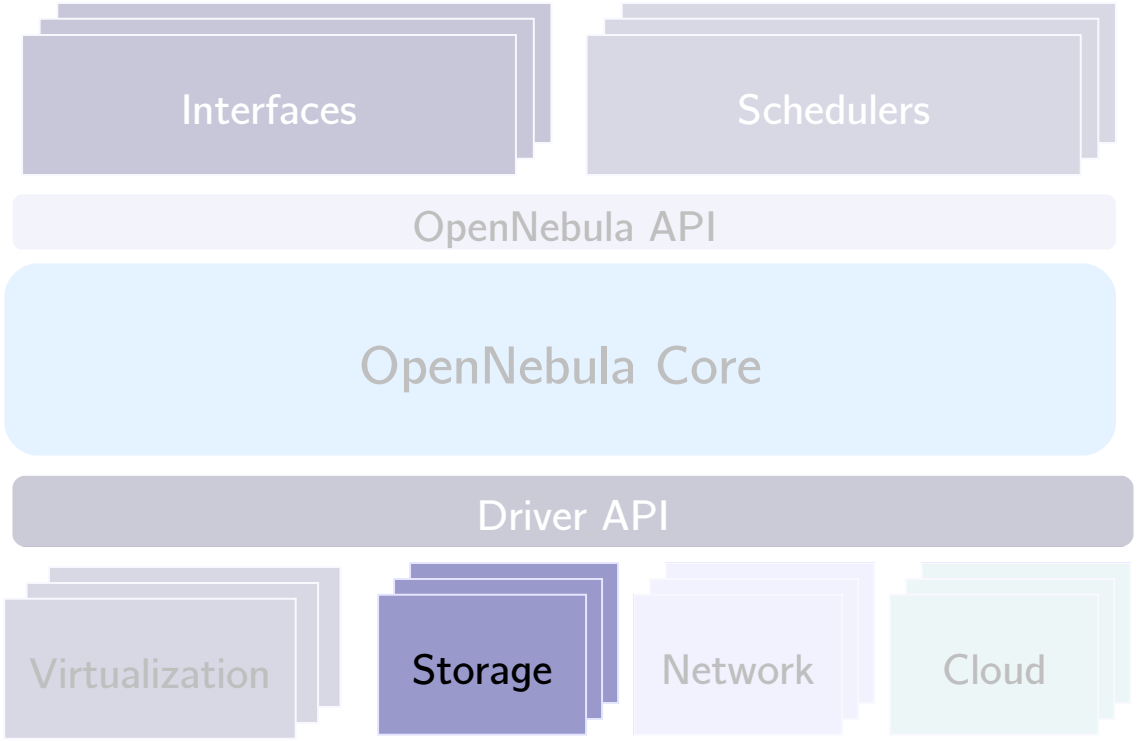

scp-wave

Disk image deployment in O(log n) time

Developed at Clemson University

https://code.google.com/p/scp-wave/

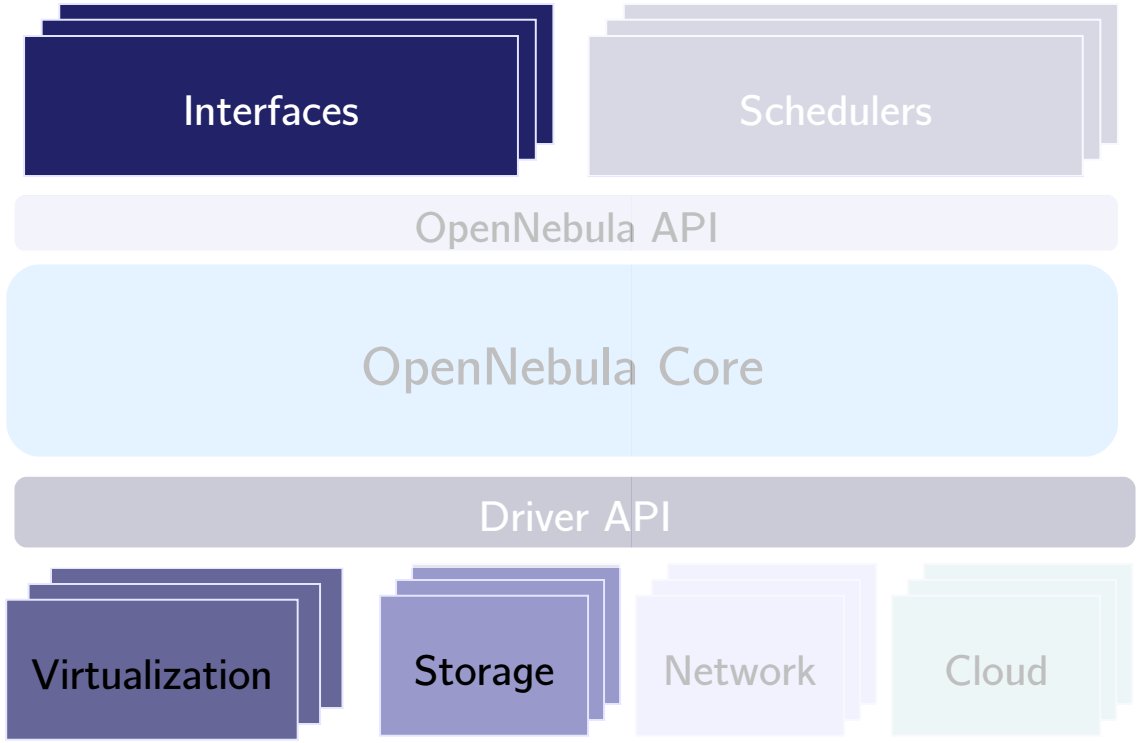

Virtual Cluster Tool

Instantiate, deploys and saves virtual clusters as atomic, self-consistent entities.

Developed at Center for Advanced Studies, Research and Development in Sardinia

http://dc.crs4.it/projects/vida

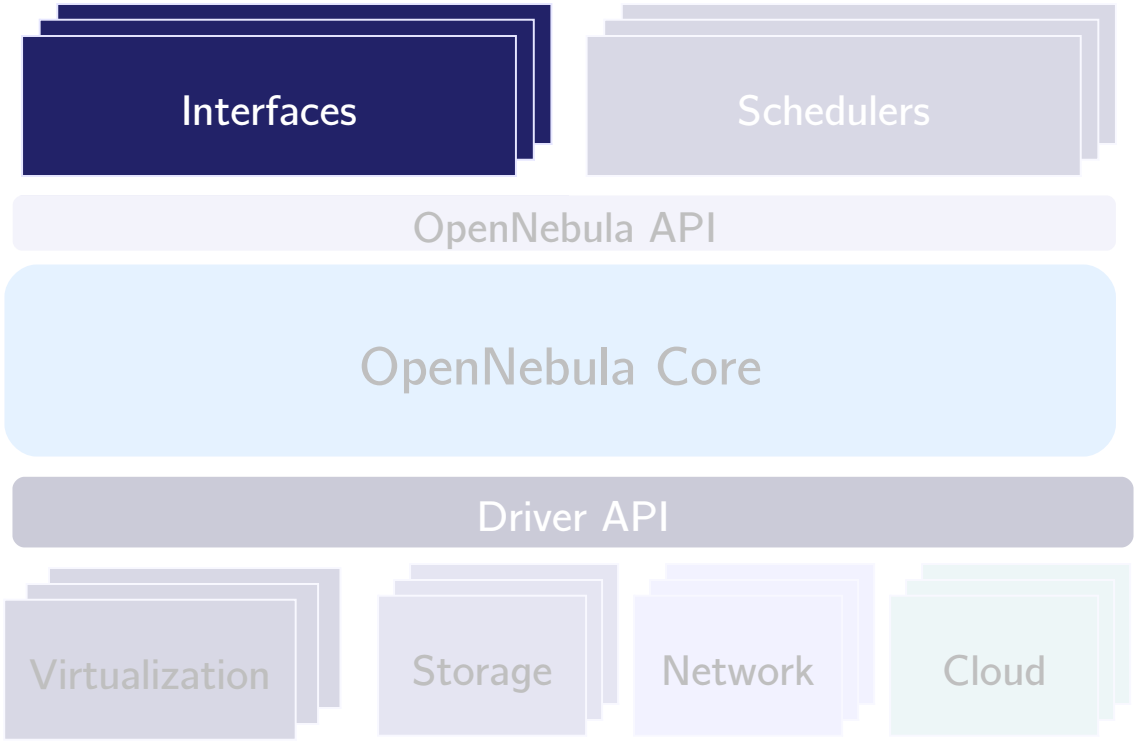

Deltacloud driver

OGF OCCI API driver for the RedHat DeltaCloud Framework.

http://deltacloud.org/drivers.html

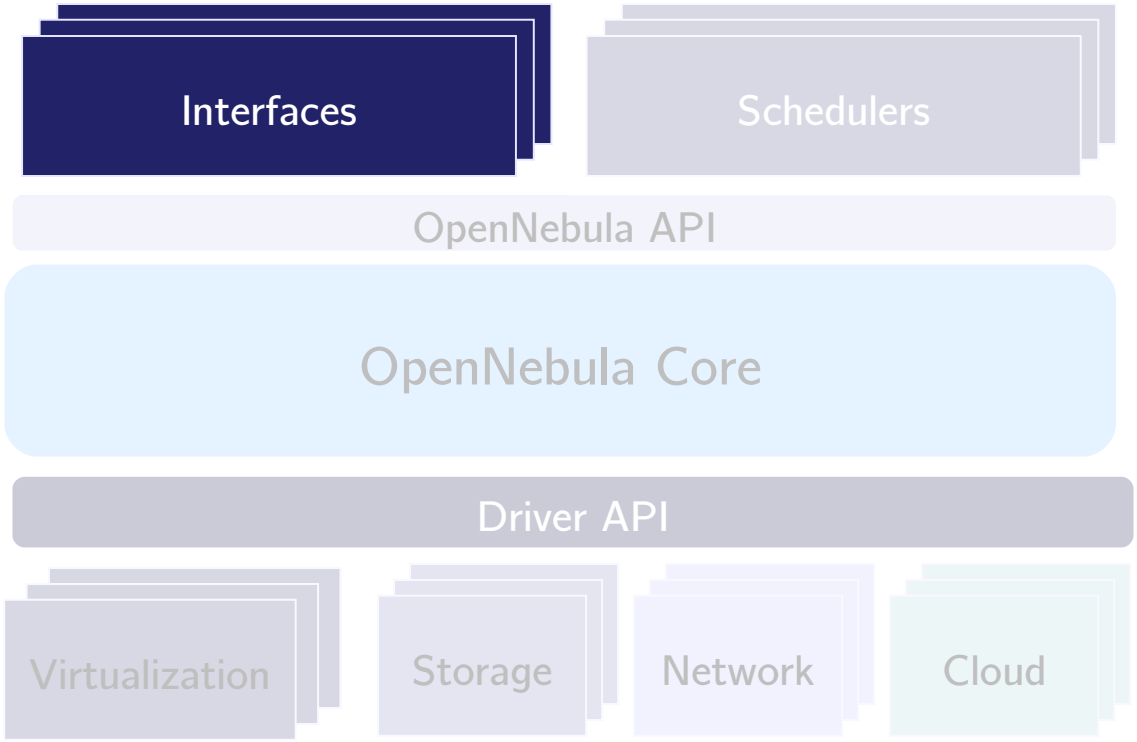

Libcloud driver

OGF OCCI API driver for Apache Libcloud

http://incubator.apache.org/libcloud/

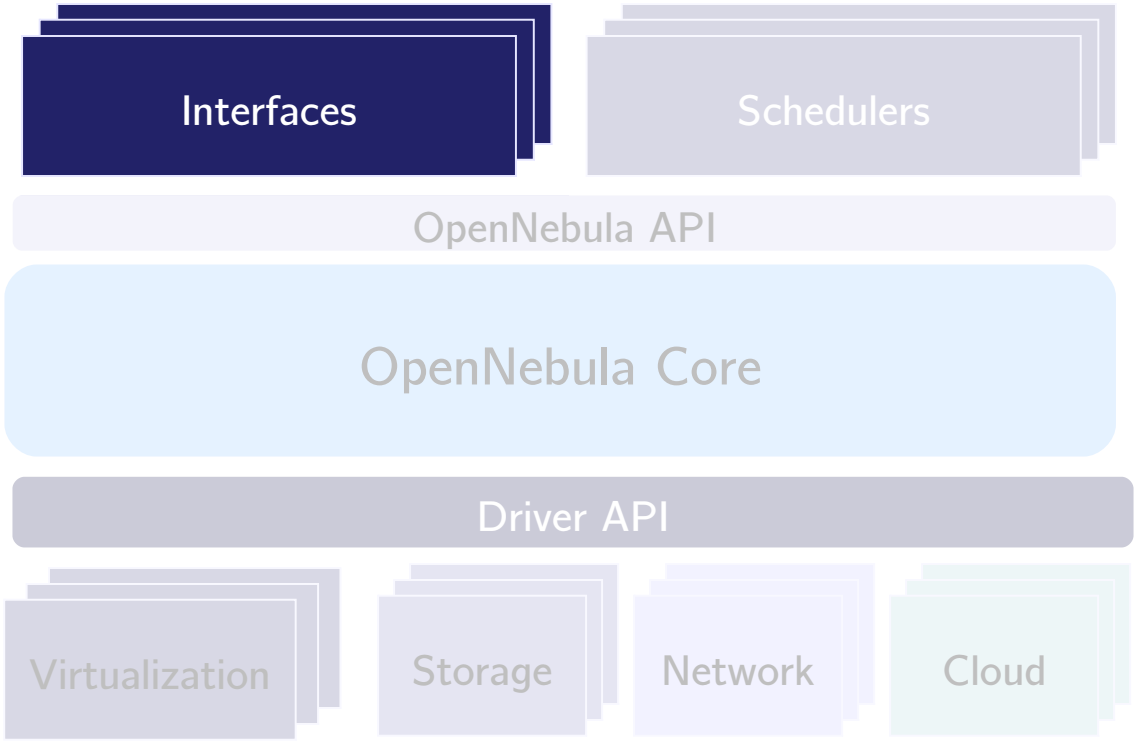

Management Console

Web interface for OpenNebula

Developed at SARA, Academic HPC center of Amsterdam

http://dev.opennebula.org/projects/management-console

Google Summer of Code 2010 projects

Management Console (II)

Virtual Environment Builder Integration

Service Manager

Improving Haizea+OpenNebula integration

The OpenNebula Cloud Toolkit: Experiences and Outlook

Borja Sotomayor
University of Chicago

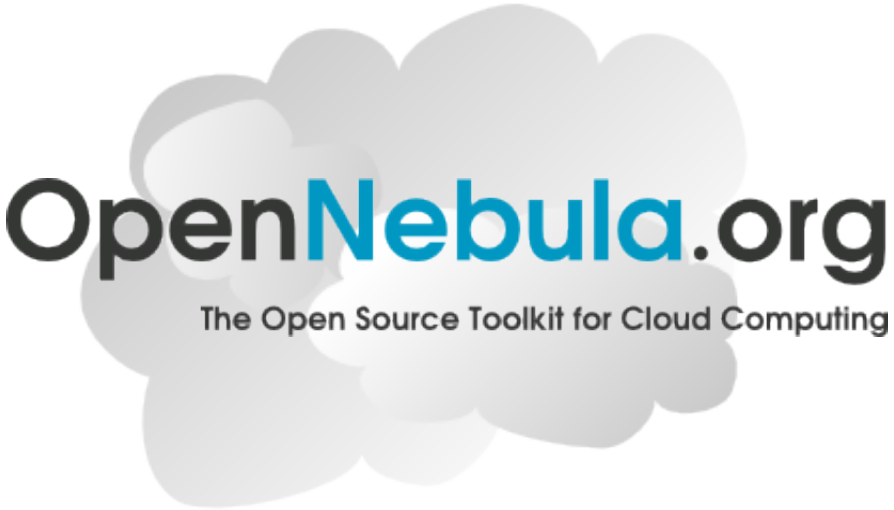

What is OpenNebula?

Experiences

Ecosystem

Outlook

On the road to OpenNebula 1.6

Addressing bottlenecks that arise when scaling to 10,000's of VMs.

Providing image repository functionality

Adding authentication and authorization drivers

Improving compatibility with EC2

And many more:
http://dev.opennebula.org/versions/show/11

Research funding guaranteed until 2013

- RESERVOIR– Resources and Services Virtualization without Barriers, 2008-2011, EU grant agreement 215605
- HPCcloud - Distributed Virtual Infrastructures to Provision Resources, 2010-2012, MICINN TIN2009-07146
- NUBA - Normalized Usage of Business-oriented Architectures, 2009-2011, MITyC Avanza TSI-020301-2009-30
- MEADIANET - Integración de Servicios Multimedia de Siguiente Generación en la Internet del Futuro, 2010-2013, Comunidad de Madrid S2009/TIC-1468
- Recently approved: StratusLab, BonFIRE, 4CAAST

StratusLab www.stratuslab.org

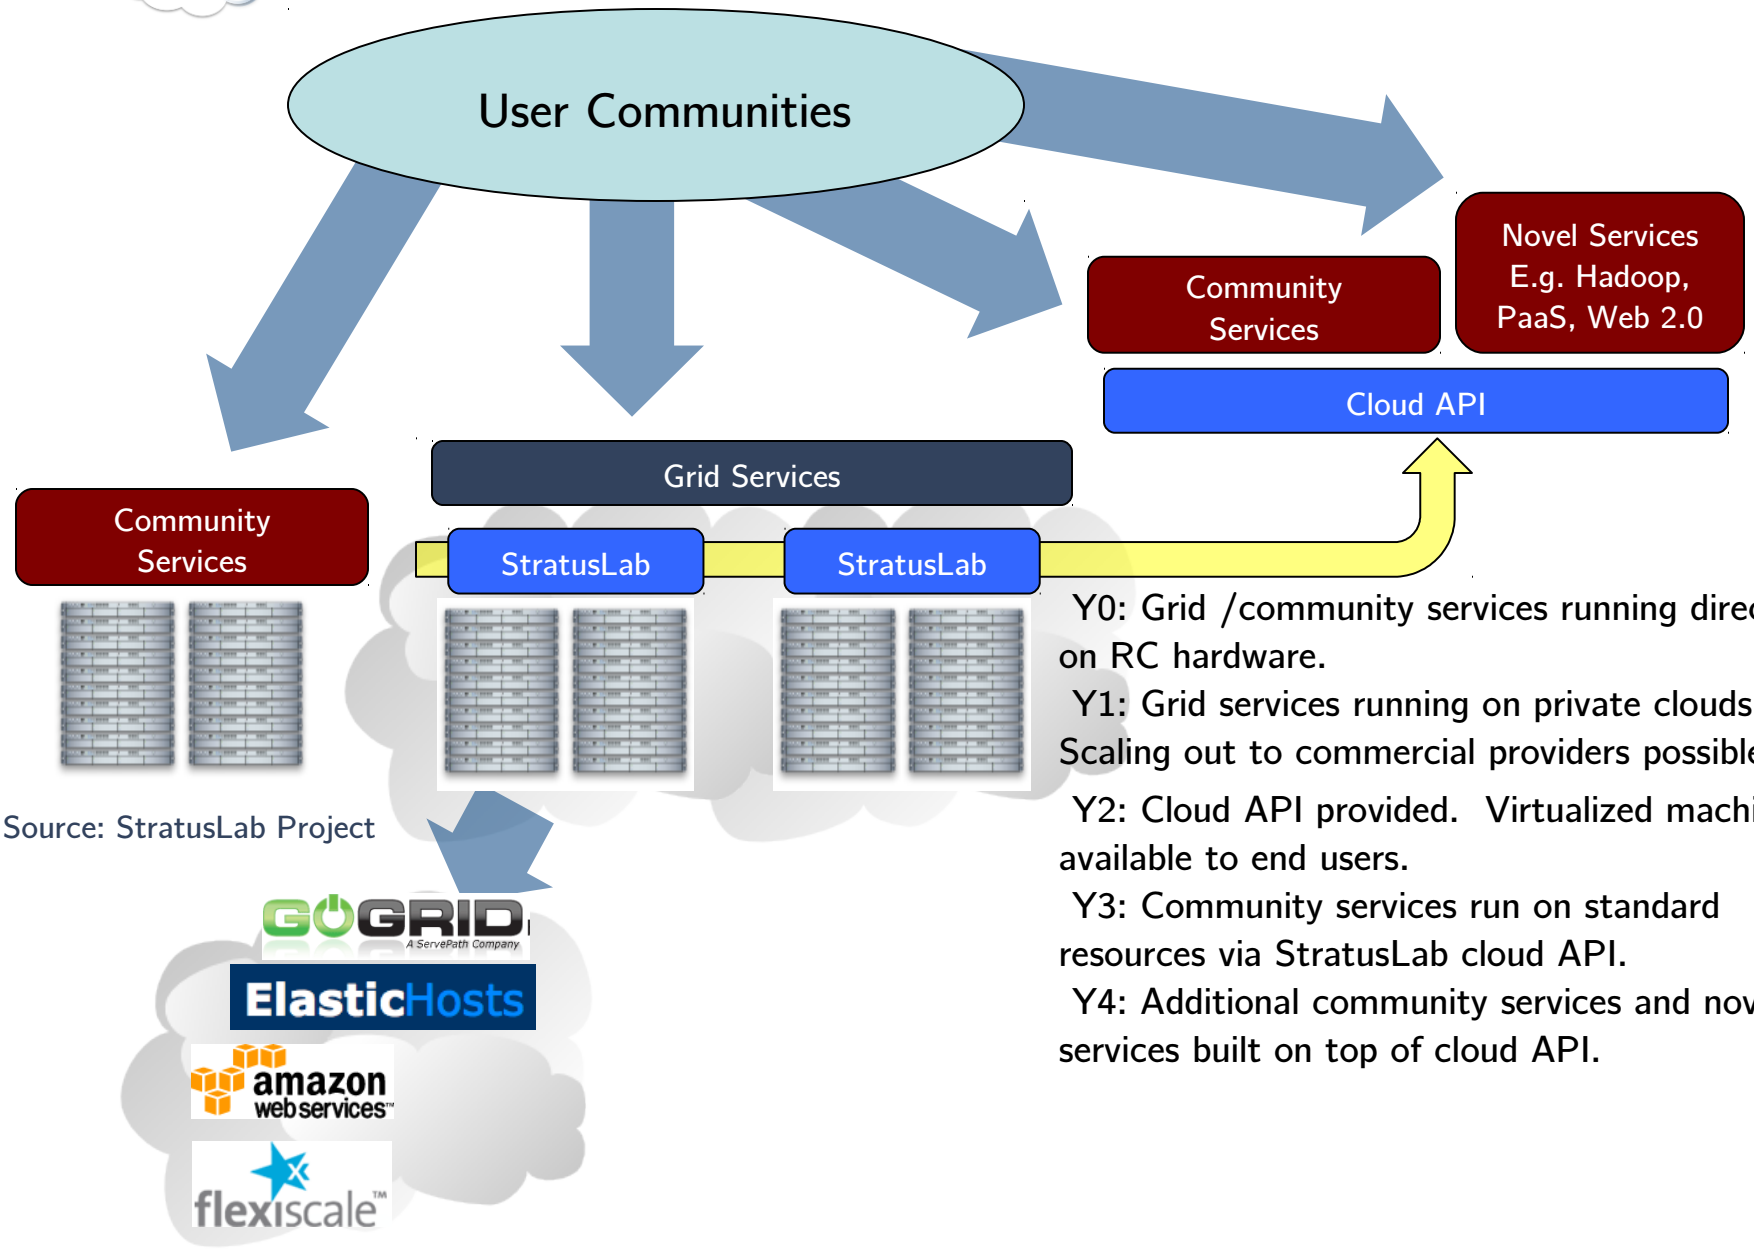

Source: StratusLab Project

Y0: Grid /community services running directly on RC hardware.

Y1: Grid services running on private clouds. Scaling out to commercial providers possible.

Y2: Cloud API provided. Virtualized machines available to end users.

Y3: Community services run on standard resources via StratusLab cloud API.

Y4: Additional community services and novel services built on top of cloud API.

Building Service Testbeds on FIRE

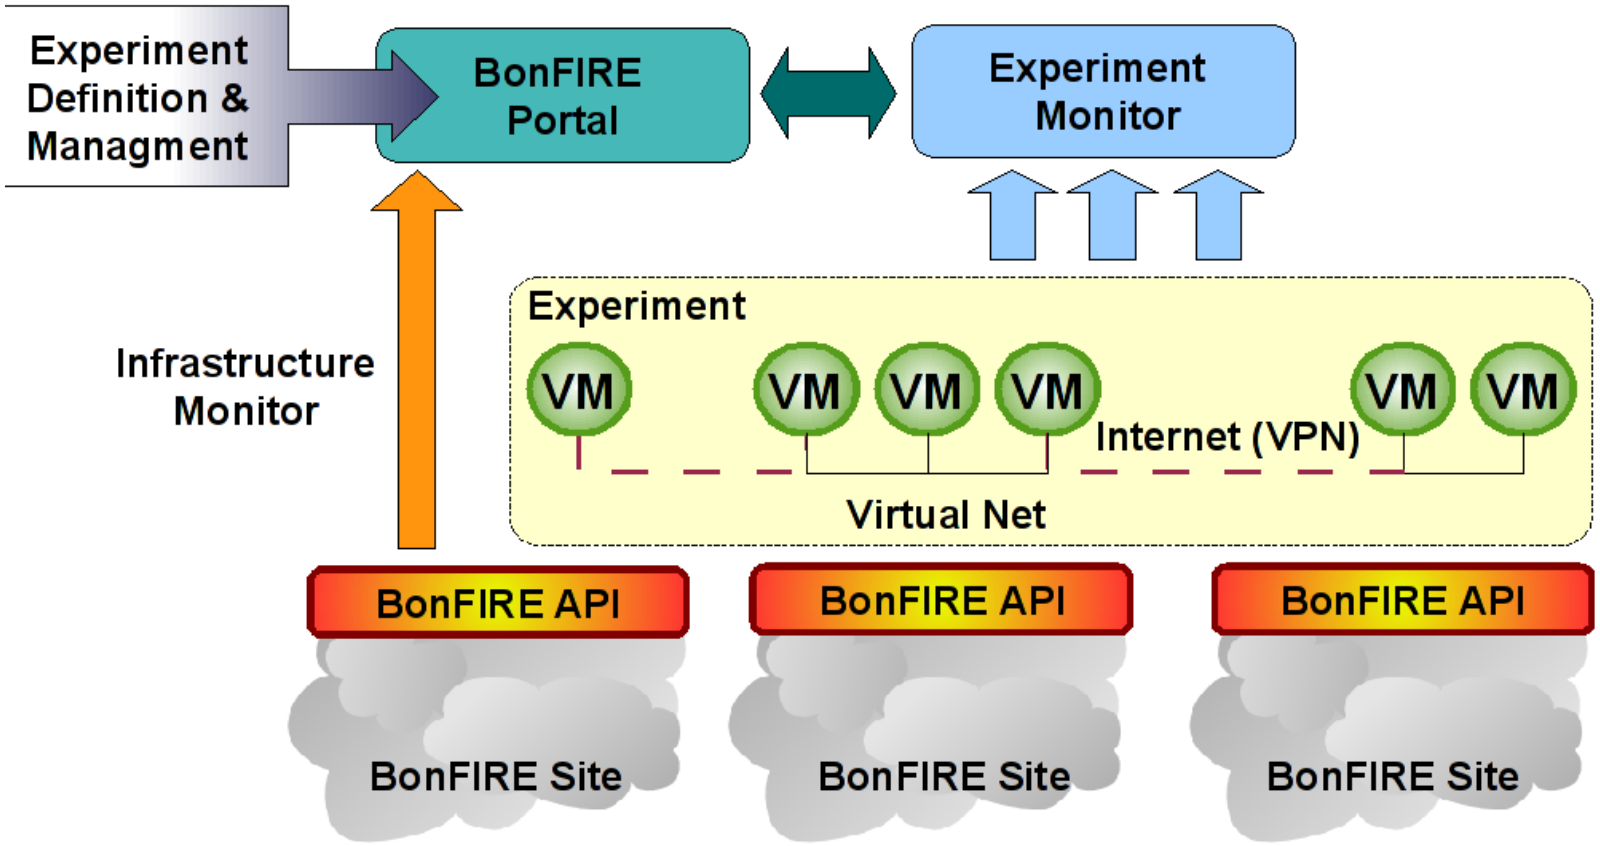

Source: BonFIRE Project

Part of European Future Internet strategy: http://www.future-internet.eu/

C12G LABS

C12G Labs is a privately-held, self-funded company, started by the authors of OpenNebula in order to provide value-added enterprise-grade solutions around it.

http://www.c12g.com/

Questions?

The OpenNebula Cloud Toolkit: Experiences and Outlook

Borja Sotomayor
University of Chicago
borja@cs.uchicago.edu

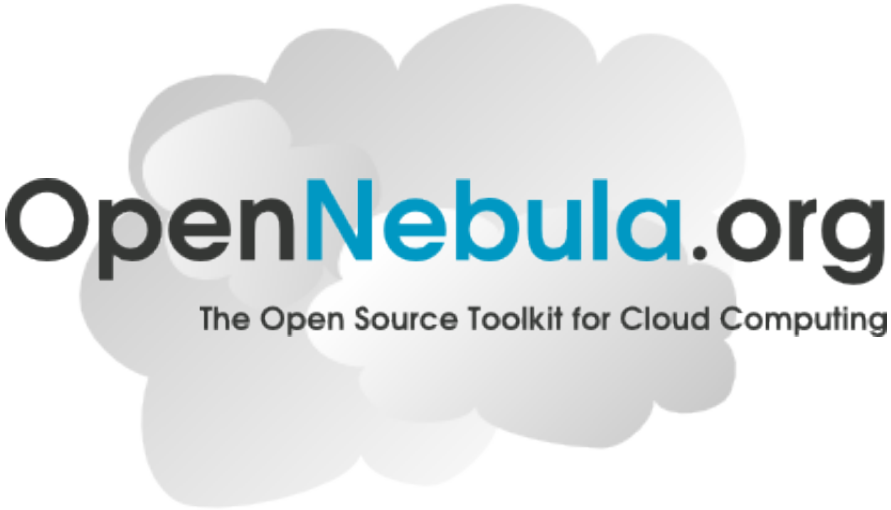

Backup Slides

Model	Definition	Examples of Deployment
Private	Infrastructure is owned by a single organization and made available only to the organization	• Optimize and simplify **internal operation** • **SaaS**/**PaaS** support • IT consolidation within **large organizations** (Goverment Clouds, University Clouds…)
Public	Infrastructure is owned by a single organization and made available to other organizations	• **Commercial cloud providers** • **Community public clouds** by ICT service centers to enable scientific and educational projects to experiment with cloud computing • **Special purpose clouds** with dedicated capabilities (Science Clouds, HPC Clouds..) • **Regional clouds** to address regulatory or latency issues
Hybrid	Infrastructure is a composition of two or more clouds	• **Cloudbursting** to address peak demands • **Cloud Federation** to share infrastructure with partners • **Cloud Aggregation** to provide a larger resource infrastructure

	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
Flexibility and Extensibility	Yes	No	Yes	Yes	Yes
Open Source	No	No	GPL	Apache	Apache