2nd June 2010
CERN Seminar
Geneva, Switzerland

Building Clouds with OpenNebula and its Application to Grid Computing

Ignacio M. Llorente

dsa-research.org

Distributed Systems Architecture Research Group
Universidad Complutense de Madrid

Position in the Cloud Ecosystem

Building Clouds with OpenNebula and its Application to Grid Computing

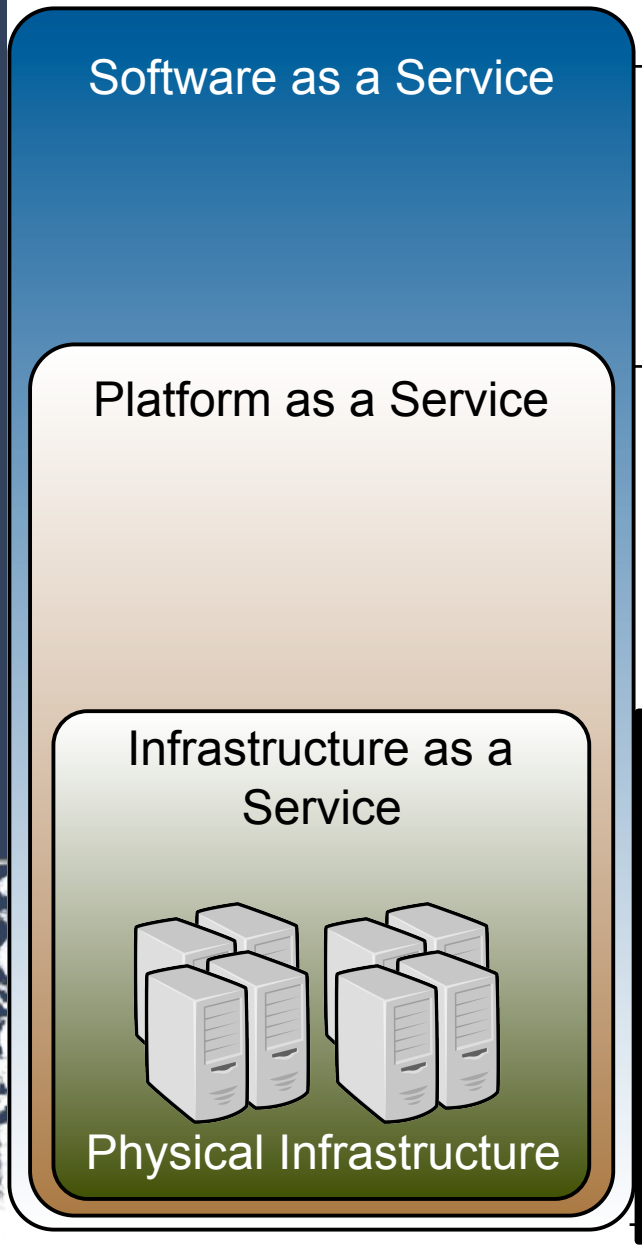

What	Who
On-demand access to any application	End-user **(does not care about hw or sw)**
Platform for building and delivering web applications	Developer **(no managing of the underlying hw & swlayers)**

OpenNebula.org

Innovative open, flexible and scalable technology to configure your own IT resources into a IaaS cloud

Transforming your IT Infrastructure into a Cloud

Building Clouds with OpenNebula and its Application to Grid Computing

Commercial Cloud Provider

- **Flexible** and **elastic capacity** to meet dynamic demands of service
- **Ubiquitous network access**
- **Pay per use** and on-demand access

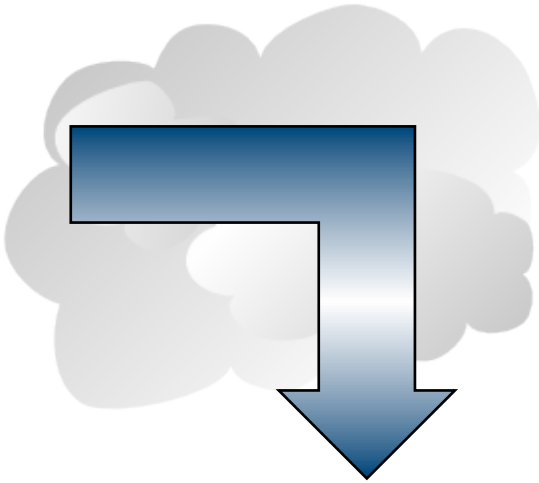

Building your Own Cloud

- **Optimize and Simplify Internal Operations**
  - **Centralized management** of all servers and services with dynamic resizing of infrastructure and dynamic allocation of capacity
  - **Higher utilization** and **operational saving** of existing resources with server consolidation and removal of application silos
  - **Lower infrastructure expenses** with combination of local and remote Cloud resources
- **Support new IT, scientific, or business Cloud services**

dsa-research.org

Deployment Models

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

dsa-research.org

Contents

Building a Cloud Infrastructure

OpenNebula as **Cloud Enabler**

A Tool to Enhance Computing Infrastructures

OpenNebula to **optimize and simplify use and operation of cluster and Grid computing infrastructures**

A Tool for Innovation

European Projects on Cloud Computing Infrastructures: **RESERVOIR, StratusLab** and **BonFIRE**

dsa-research.org

Building a Cloud: A Design Driven by Requirements

Building Clouds with OpenNebula and its Application to Grid Computing

Requirements from Usage and Deployment Scenarios

Users	Managers	Integrators	Business
Functionality exposed and workload profile	Flexible, efficient and scalable management	Open architecture, and code	Hybrid cloud computing and federation

"One solution does not fit all requirements and constraints, a properly architected solution should fully align with your IT strategy"

Constraints from Existing Infrastructure and Processes in the Organization

Building a Cloud: Flexible Cloud Manager

Building Clouds with OpenNebula and its Application to Grid Computing

Cloud Manager to Orchestrate the Complexity of a Datacenter

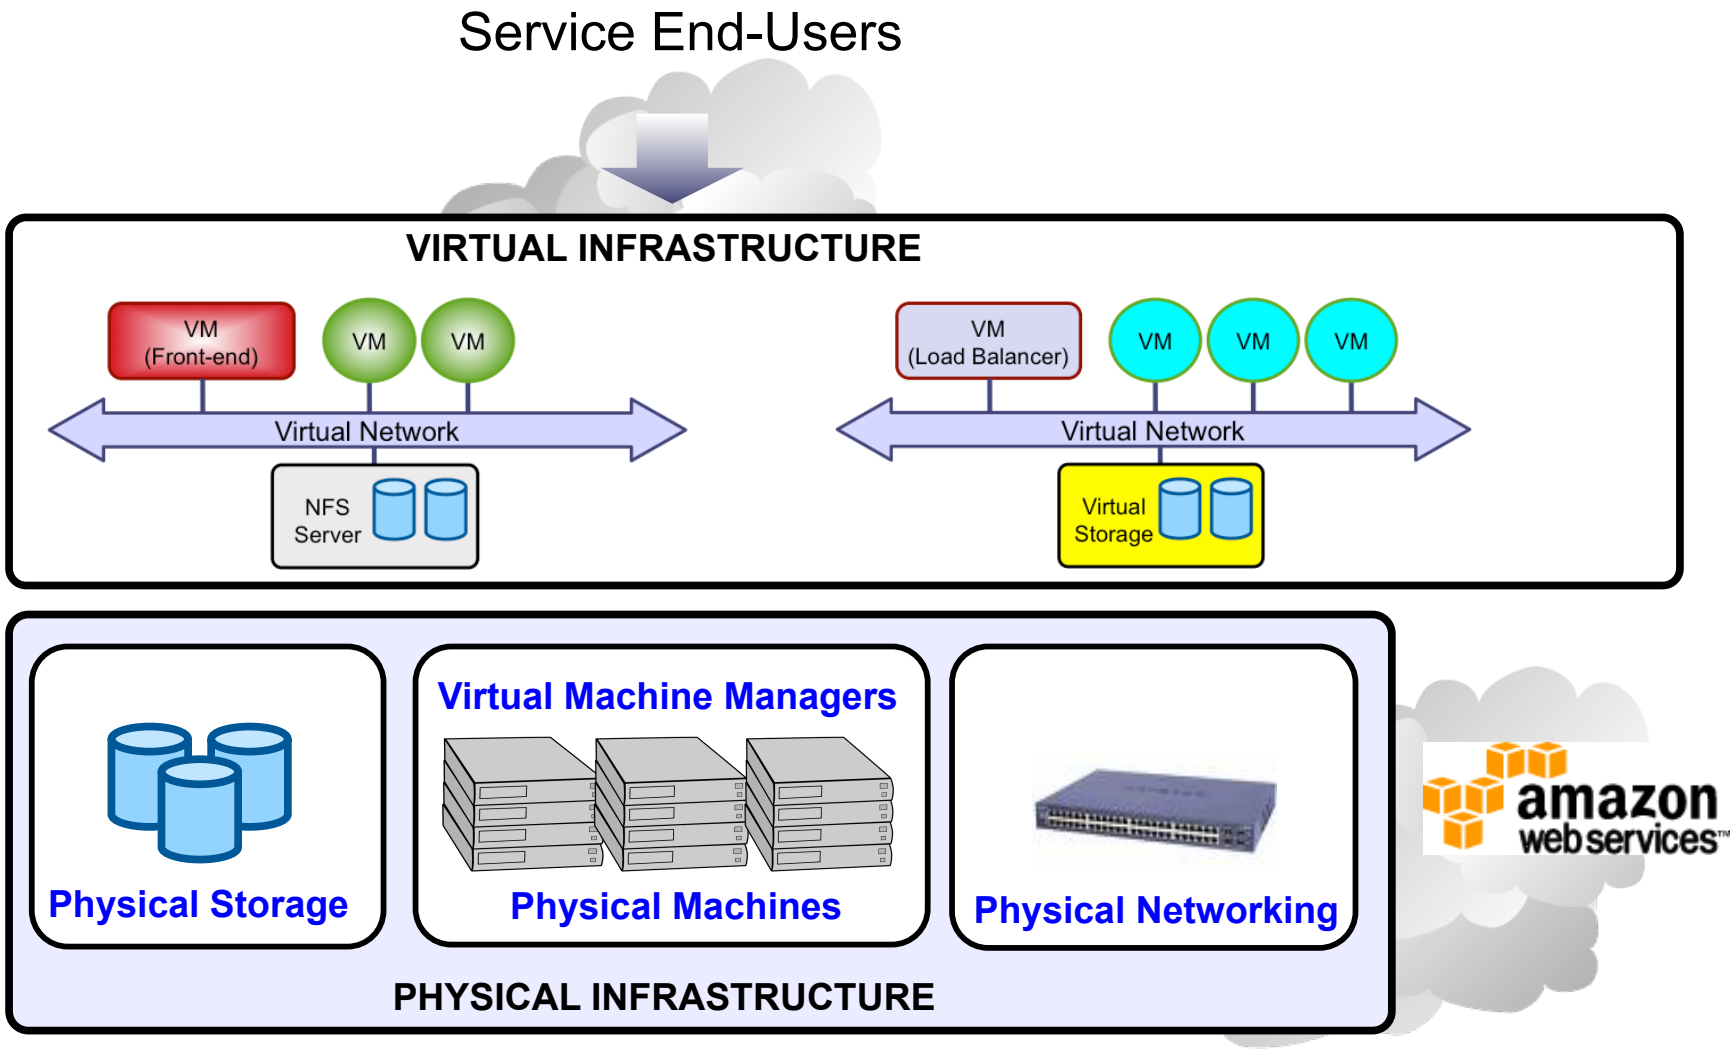

Building a Cloud: Flexible Cloud Manager

Building Clouds with OpenNebula and its Application to Grid Computing

Cloud Manager as Enabler to Build Your Own Cloud

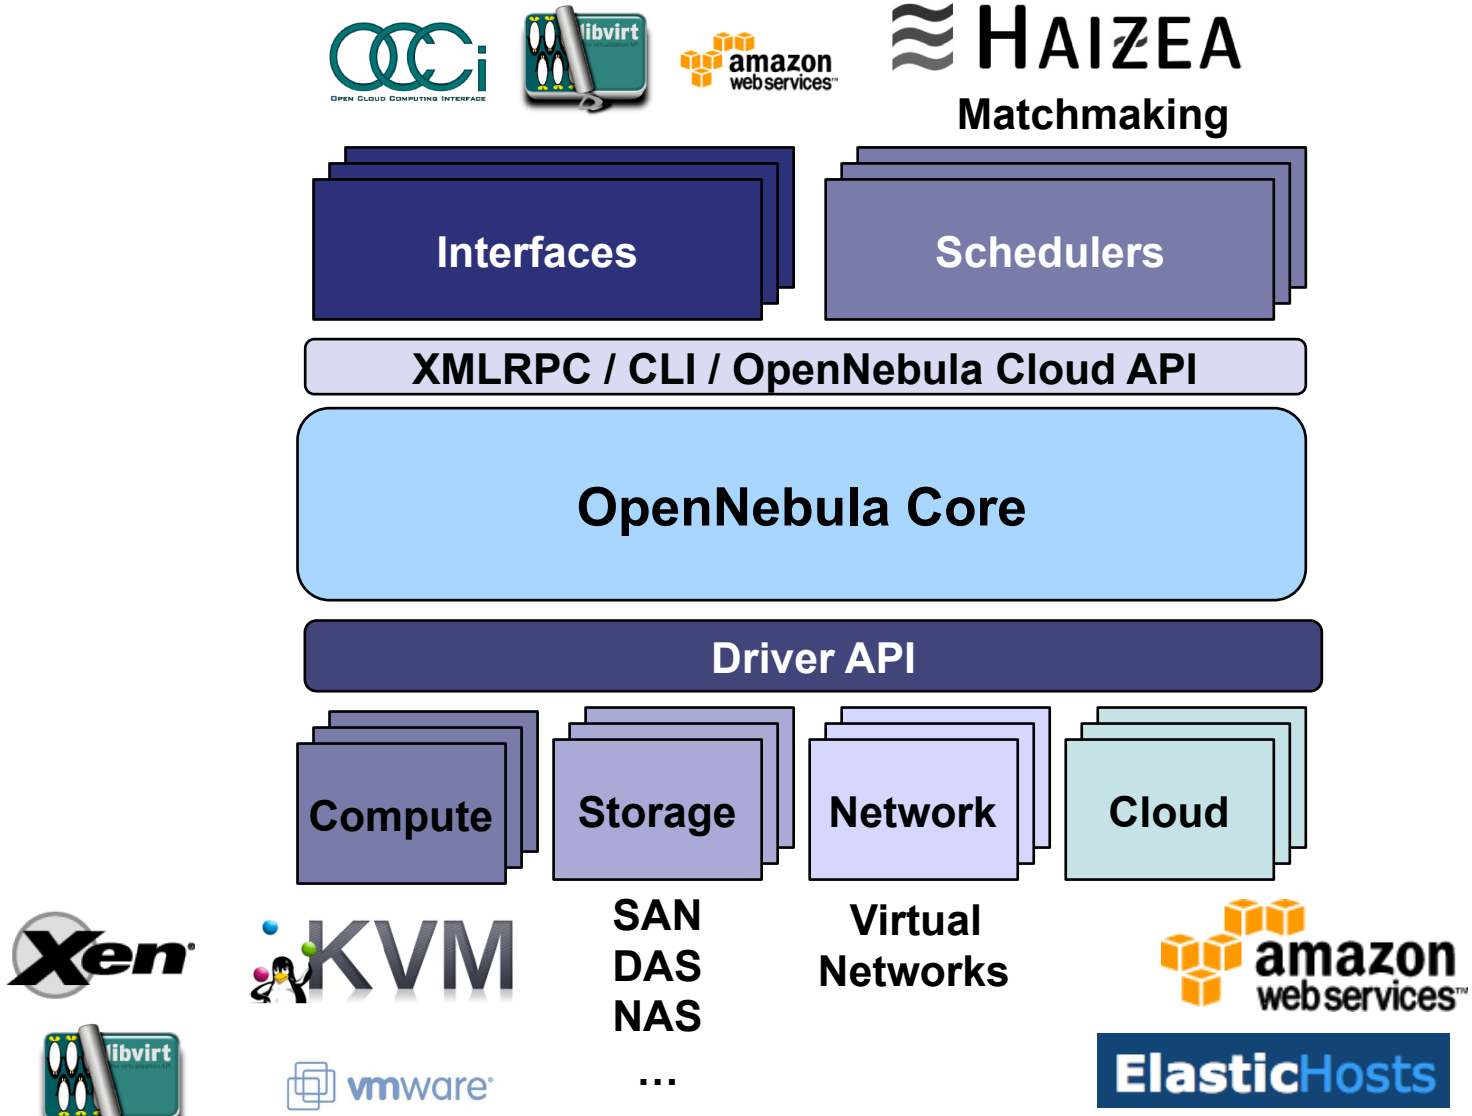

Building a Cloud: Interoperability

Building Clouds with OpenNebula and its Application to Grid Computing

Interoperation from Different Perspectives

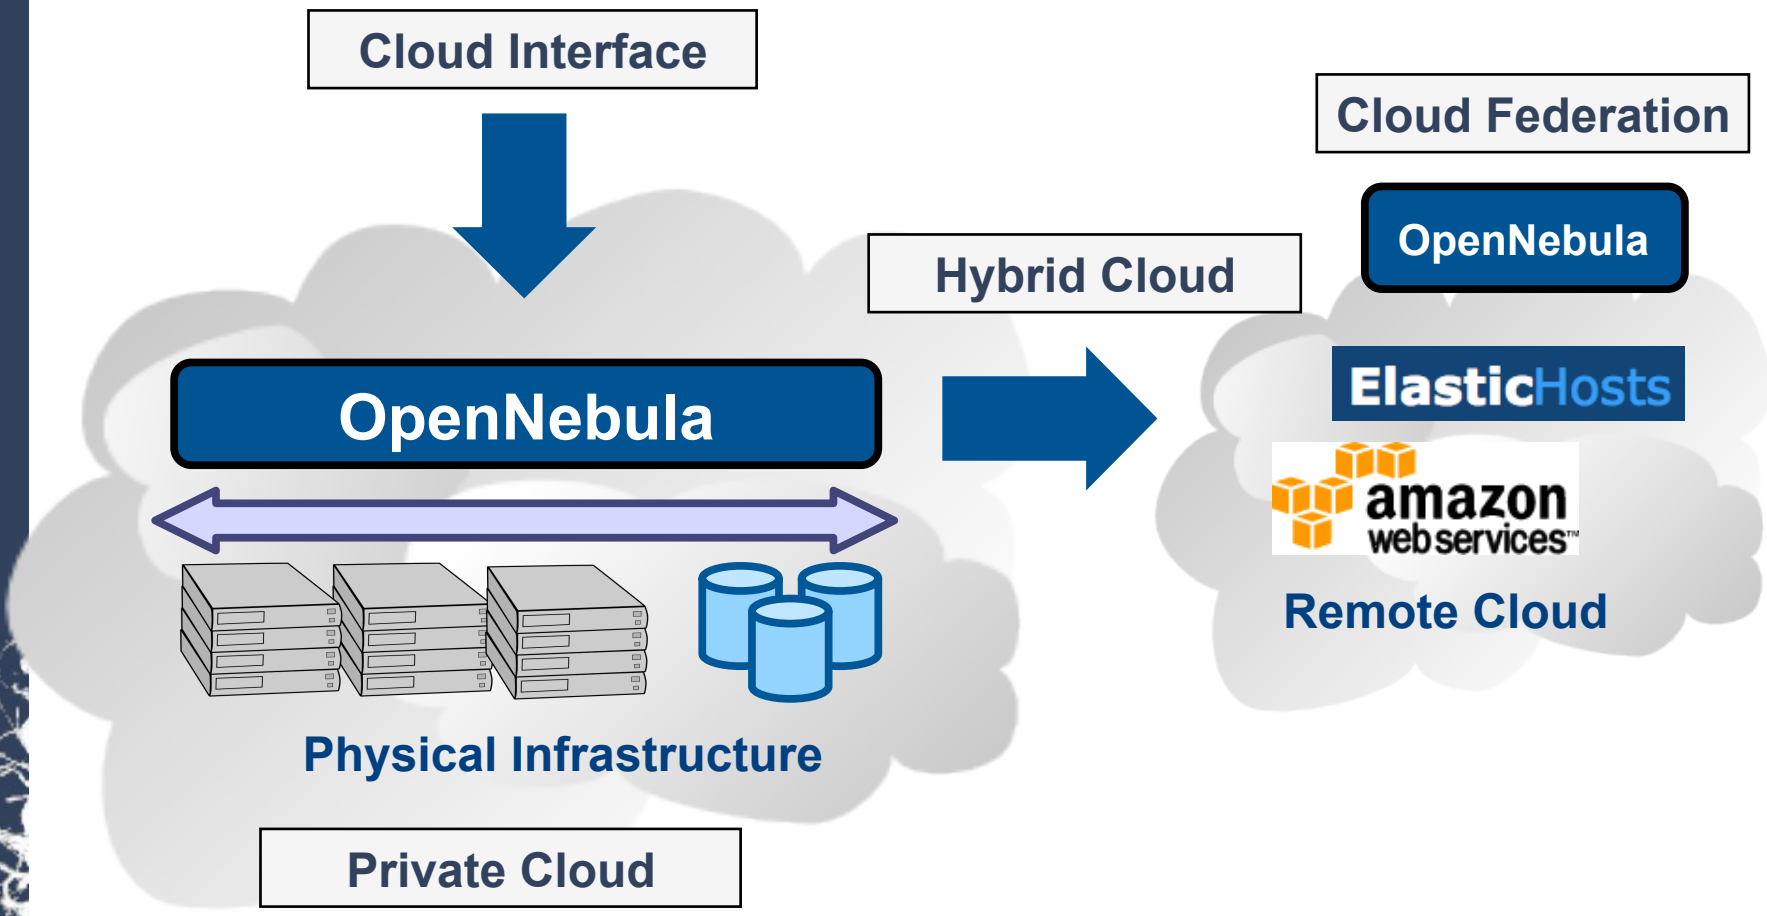

Building a Cloud: OpenNebula as Cloud Enabler

Building Clouds with OpenNebula and its Application to Grid Computing

RESERVOIR

Innovations

Technology **challenges** in cloud computing management from **business use cases**

Open-source Toolkit

OpenNebula v1.4

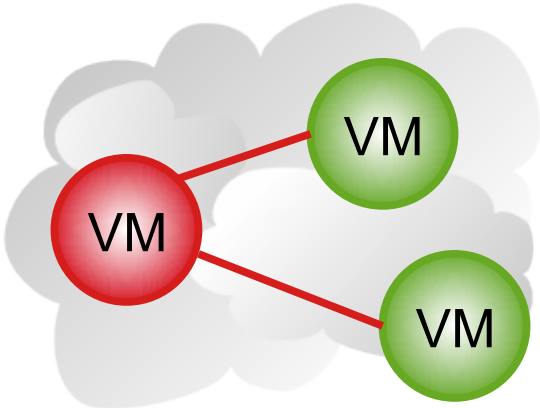

- **Open and flexible tool** to fit into any datacenter and integrate with any ecosystem component
- **Open-source** released under Apache v2.0, and distributed in Ubuntu
- **Most advanced solution** to build private, public, federated and hybrid clouds
- Based on and implements **standards to** avoid vendor lock-in and to enable interoperability
- **Efficient and scalable management** of the cloud

dsa-research.org

Building a Cloud: OpenNebula Ecosystem

Building Clouds with OpenNebula and its Application to Grid Computing

Open Community for Cloud Computing

- **Haizea Lease Manager (University of Chicago):** Advance reservation of capacity and queuing of best effort requests
- **Cloud Management Console (SARA Computing and Networking Services):** Web interface for OpenNebula
- **Virtual Cluster Tool (CRS4 Distributed Computing Group):** Atomic virtual cluster management with versioning and multiple transport protocols.
- **DeltaCloud Driver (DSA-Research@UCM)**
- **RESERVOIR Policy Engine (IBM Haifa/Elsag Datamat):** Policy-driven probabilistic admission control and dynamic placement optimization to satisfy site level management policies
- **VM Consolidation Scheduler (DSA-Research@UCM):** Periodic re-placement of VMs for server consolidation and suspension/resume of physical resources
- **Claudia (Telefonica I+D):** SLA-driven automatic service management
- **Under Development:** SUN Cloud API, vCloud API, VirtualBox plugin, dashboard for infrastructure management, new schedulers, SLA and security framework, Grid service manager, LVM and SAN support,…

A Tool to Enhance Computing Infrastructures

Building Clouds with OpenNebula and its Application to Grid Computing

Different Levels of Use: From Experimental to Production

esa esac european space astronomy centre

CERN

NIKHEF

NCHC

CESCA Centre de Supercomputació de Catalunya

CESGA

sara Computing & Networking Services

Concertation between RESERVOIR and EGEE

- Evaluate OpenNebula in the Dynamic Provisioning of EGEE Site Worker Nodes

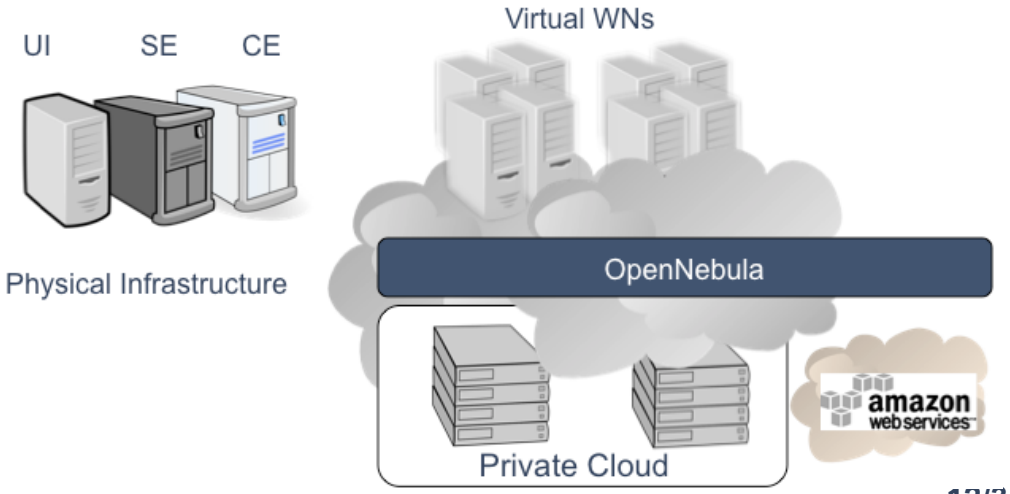

A Tool to Enhance Computing Infrastructures

Building Clouds with OpenNebula and its Application to Grid Computing

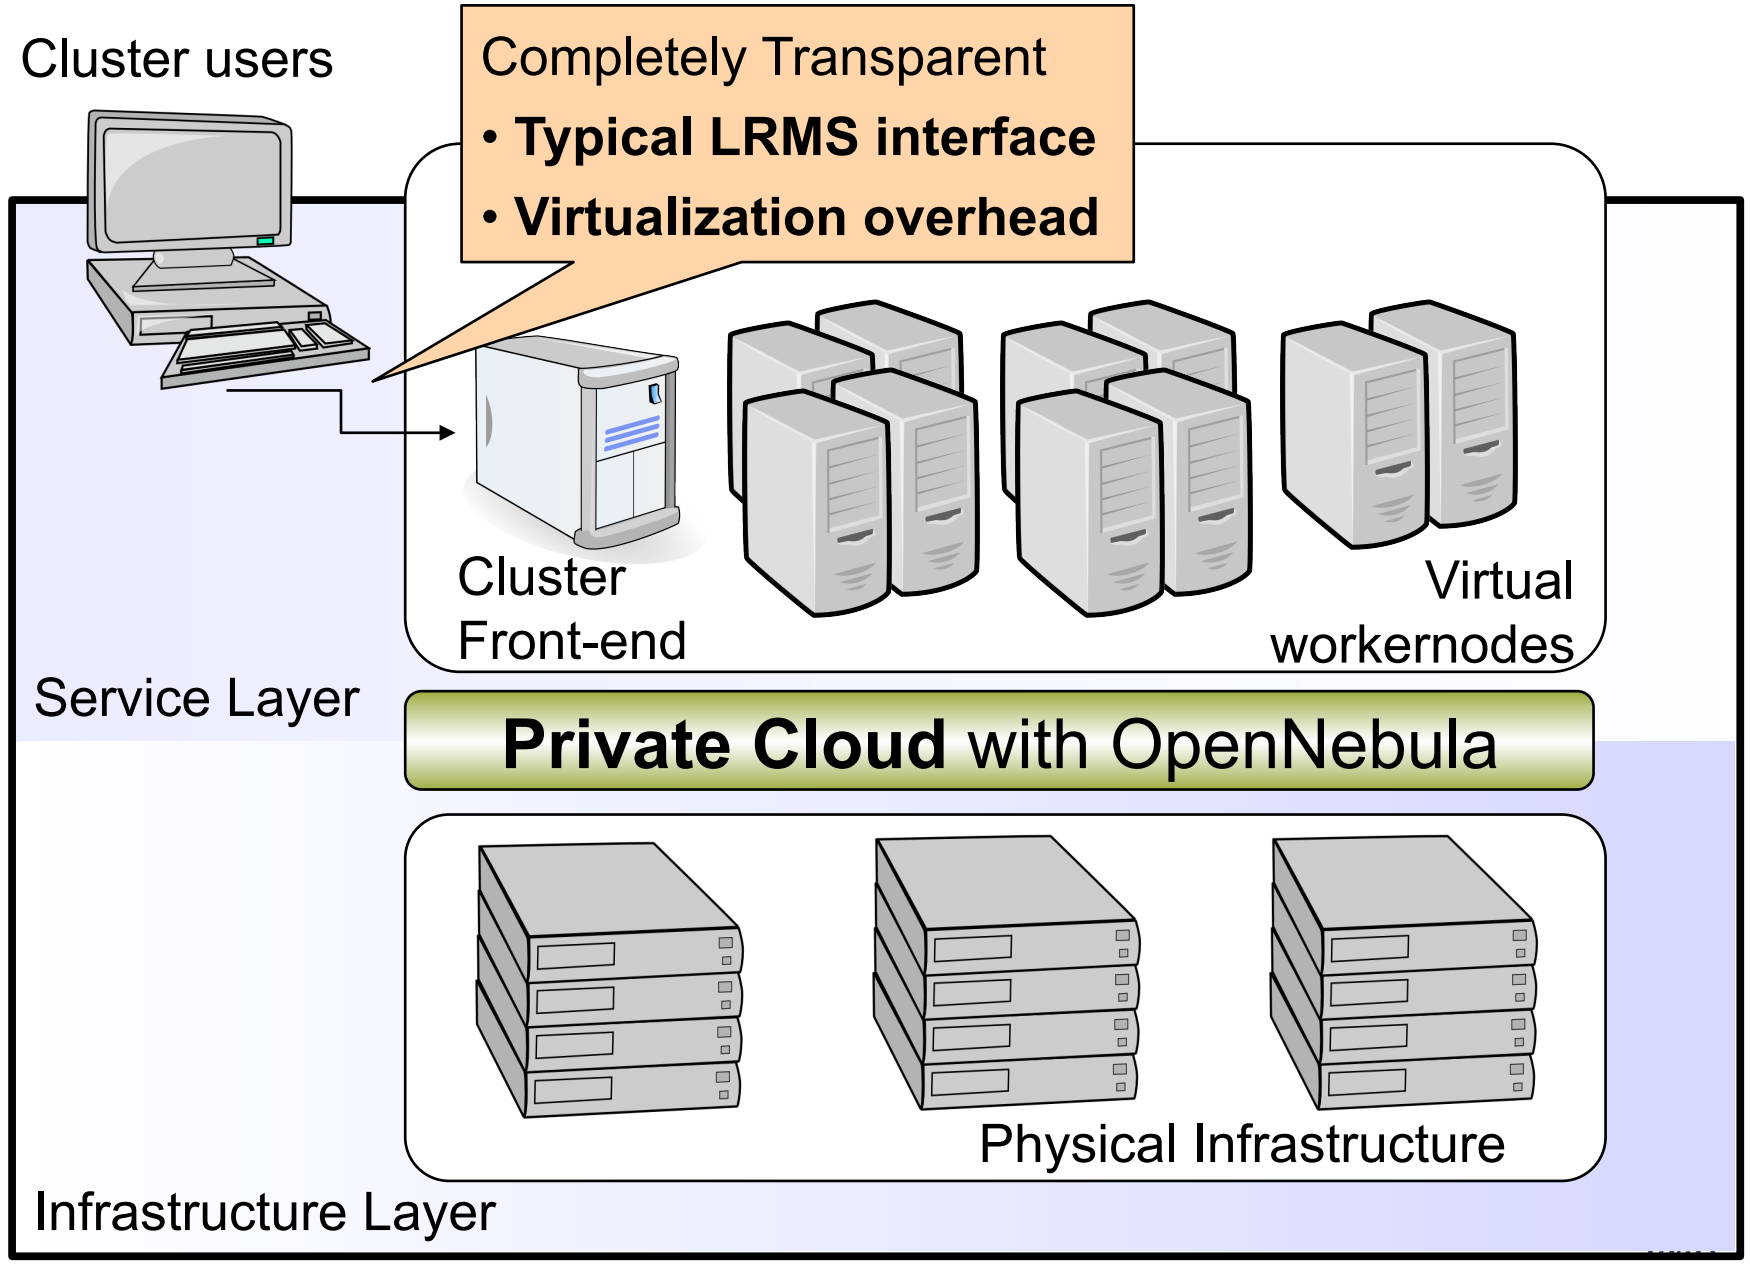

A Tool to Enhance Computing Infrastructures

Building Clouds with OpenNebula and its Application to Grid Computing

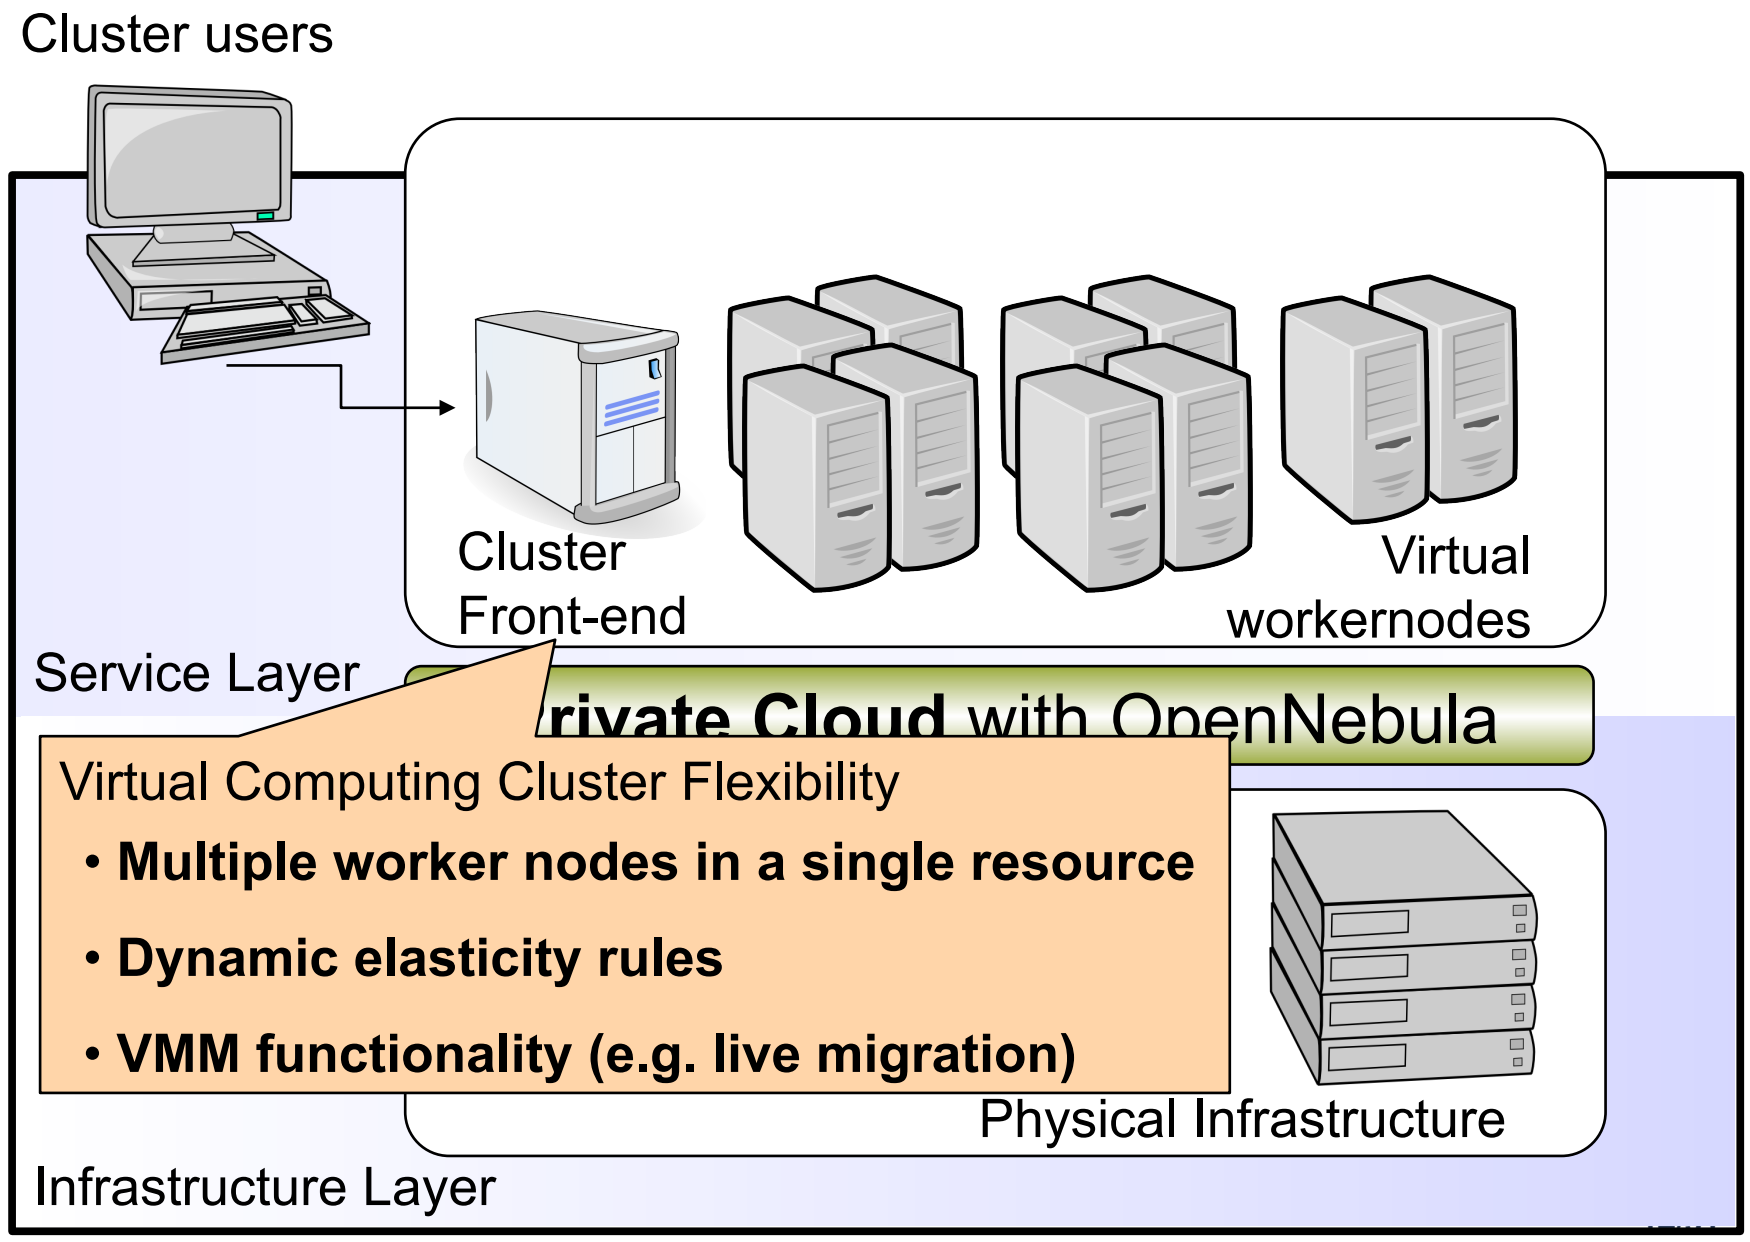

A Tool to Enhance Computing Infrastructures

Building Clouds with OpenNebula and its Application to Grid Computing

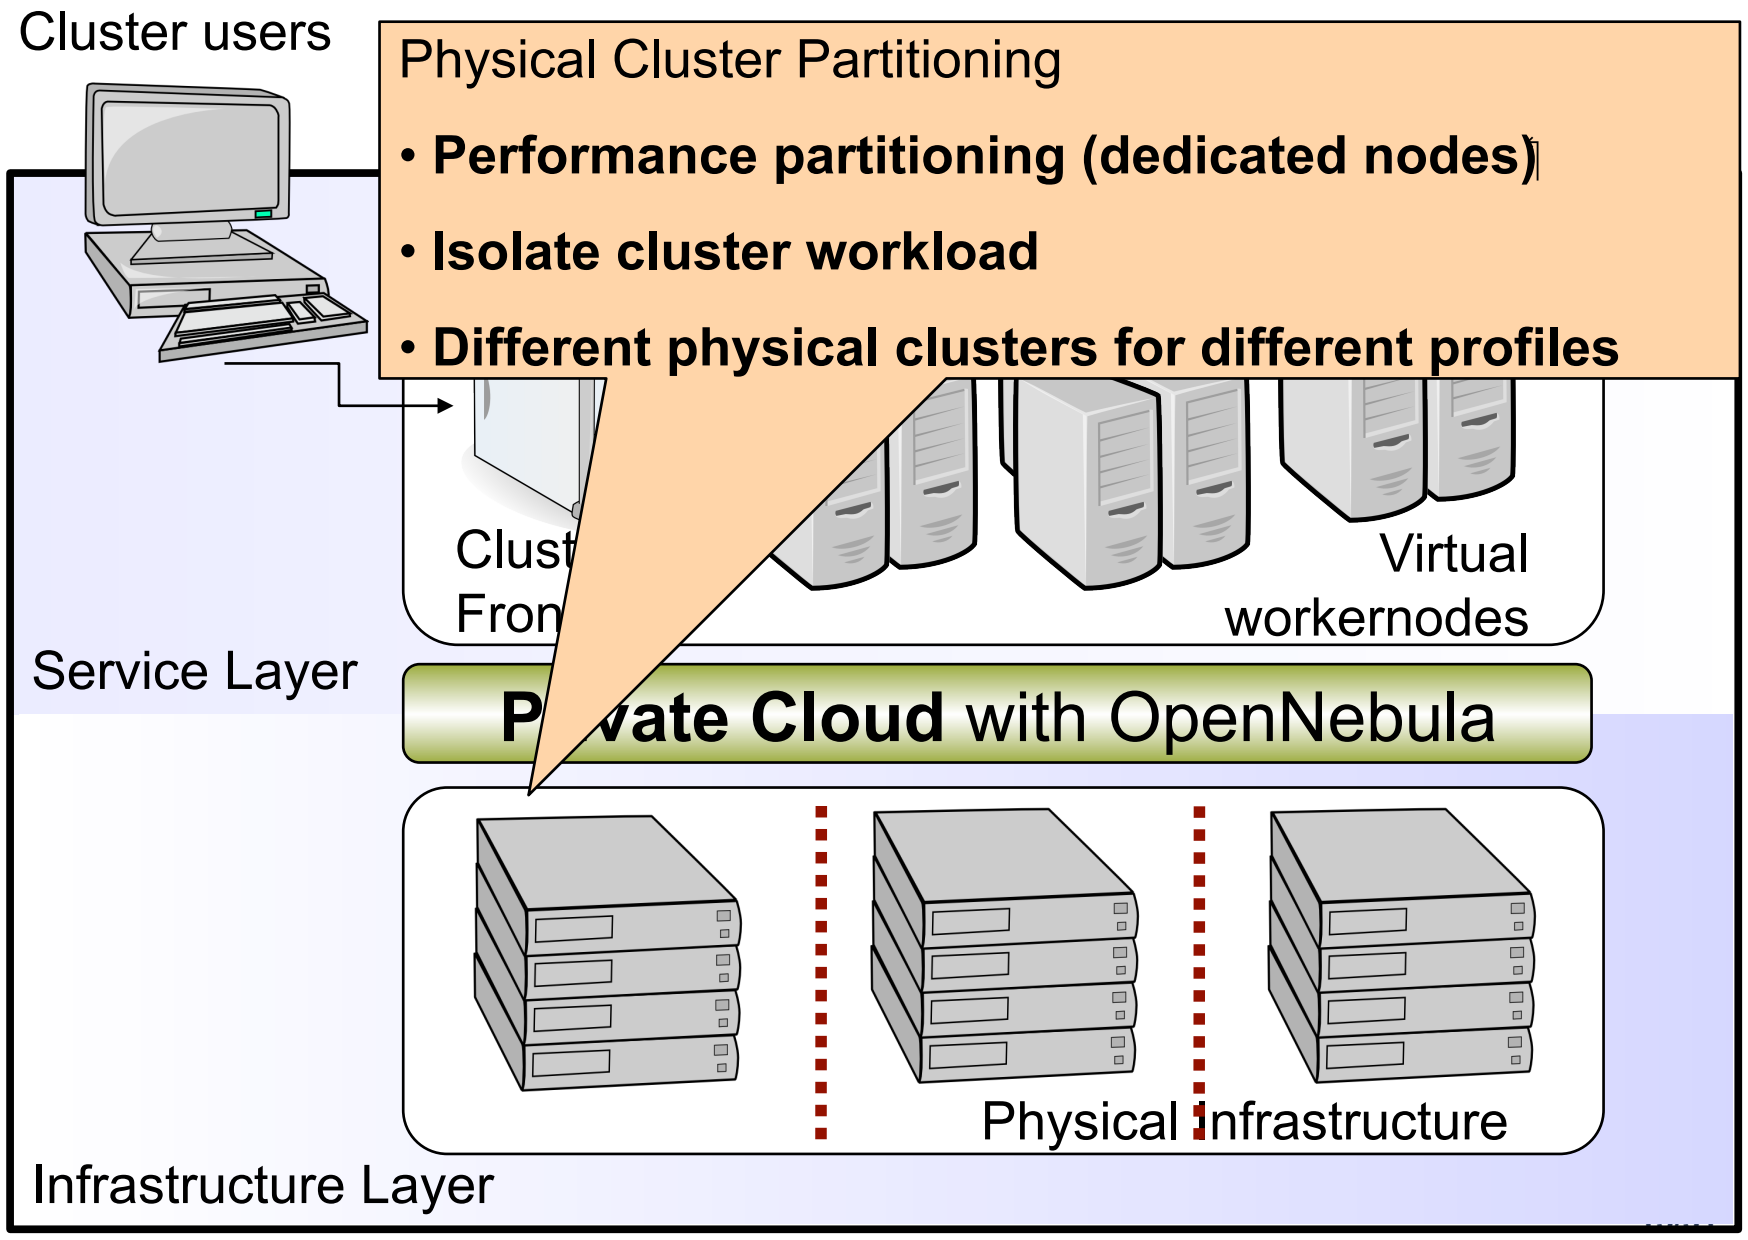

A Tool to Enhance Computing Infrastructures

Building Clouds with OpenNebula and its Application to Grid Computing

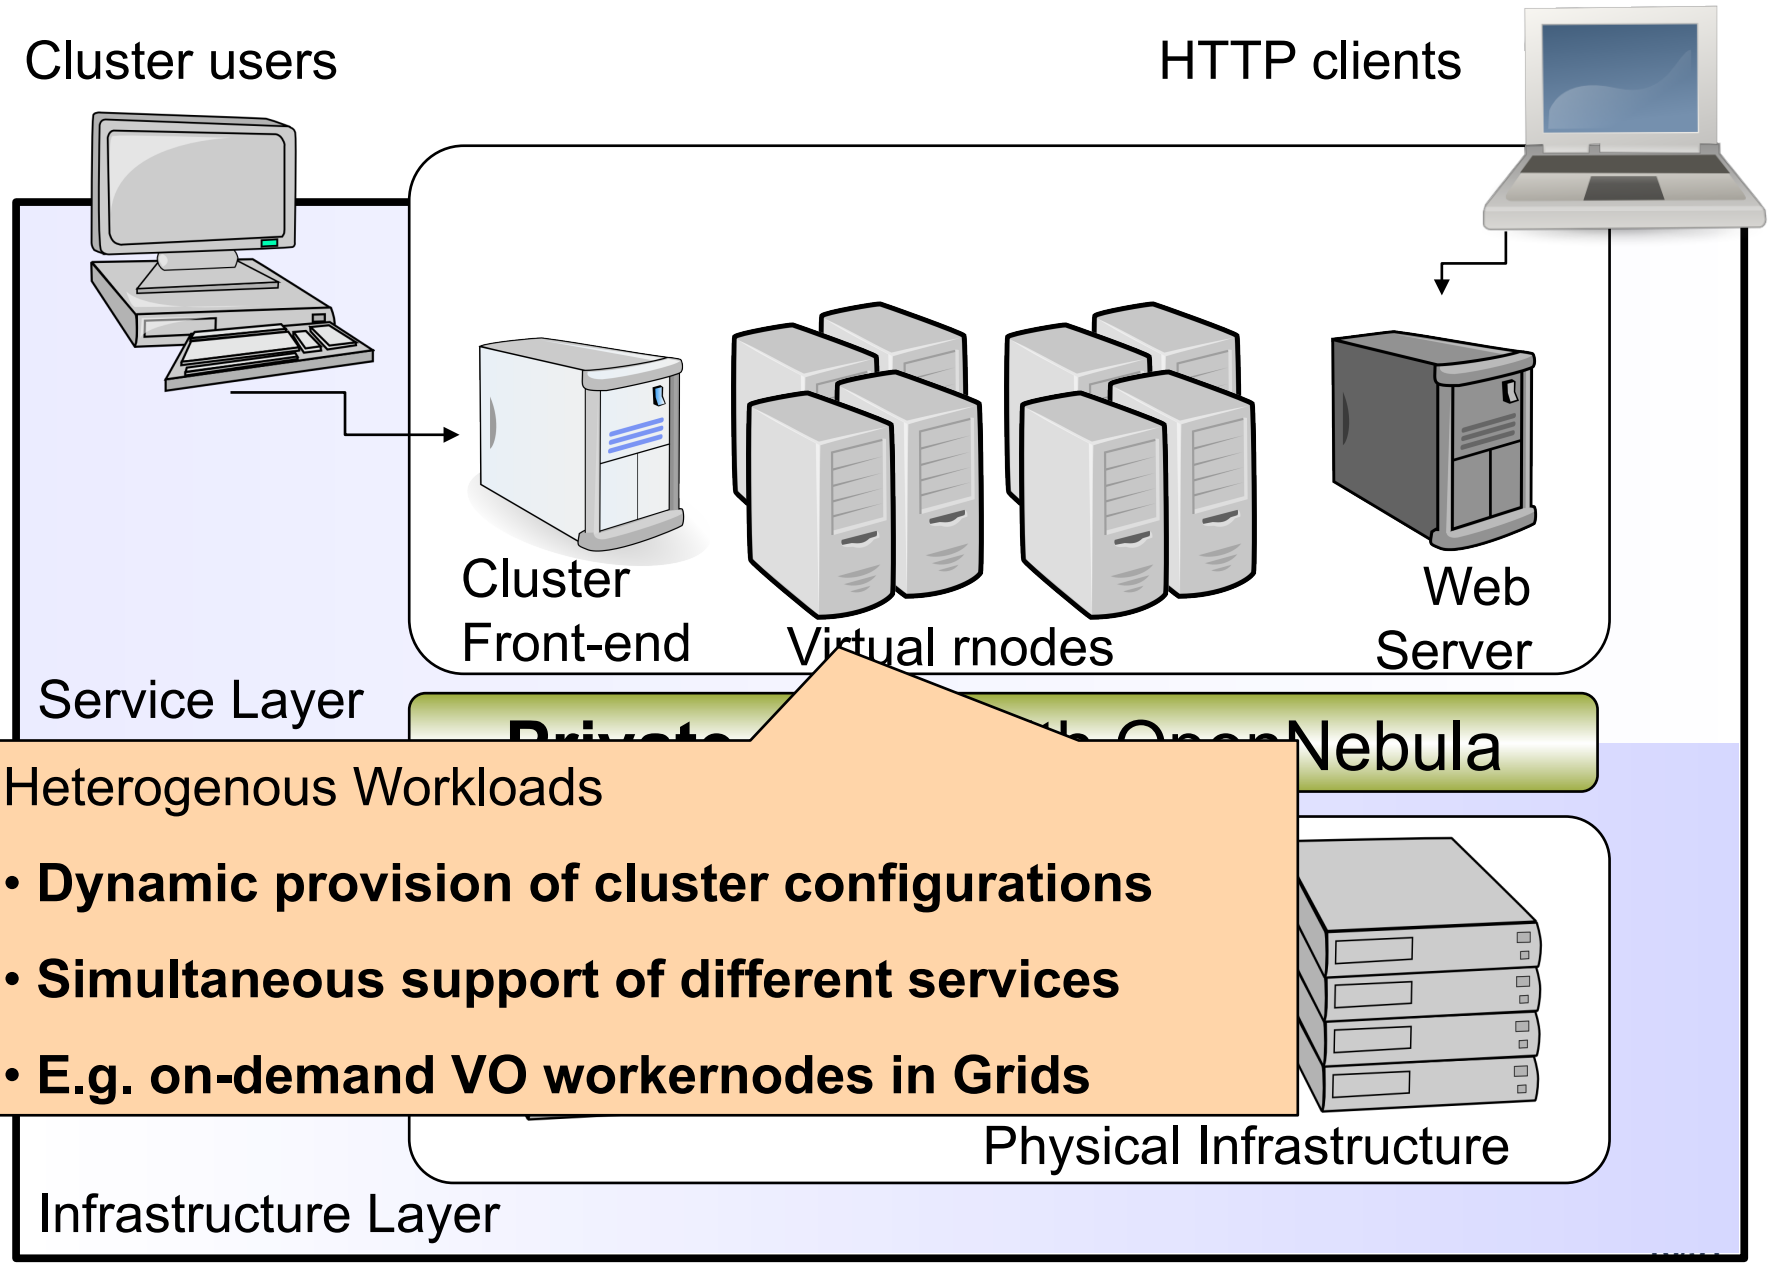

A Tool to Enhance Computing Infrastructures

Building Clouds with OpenNebula and its Application to Grid Computing

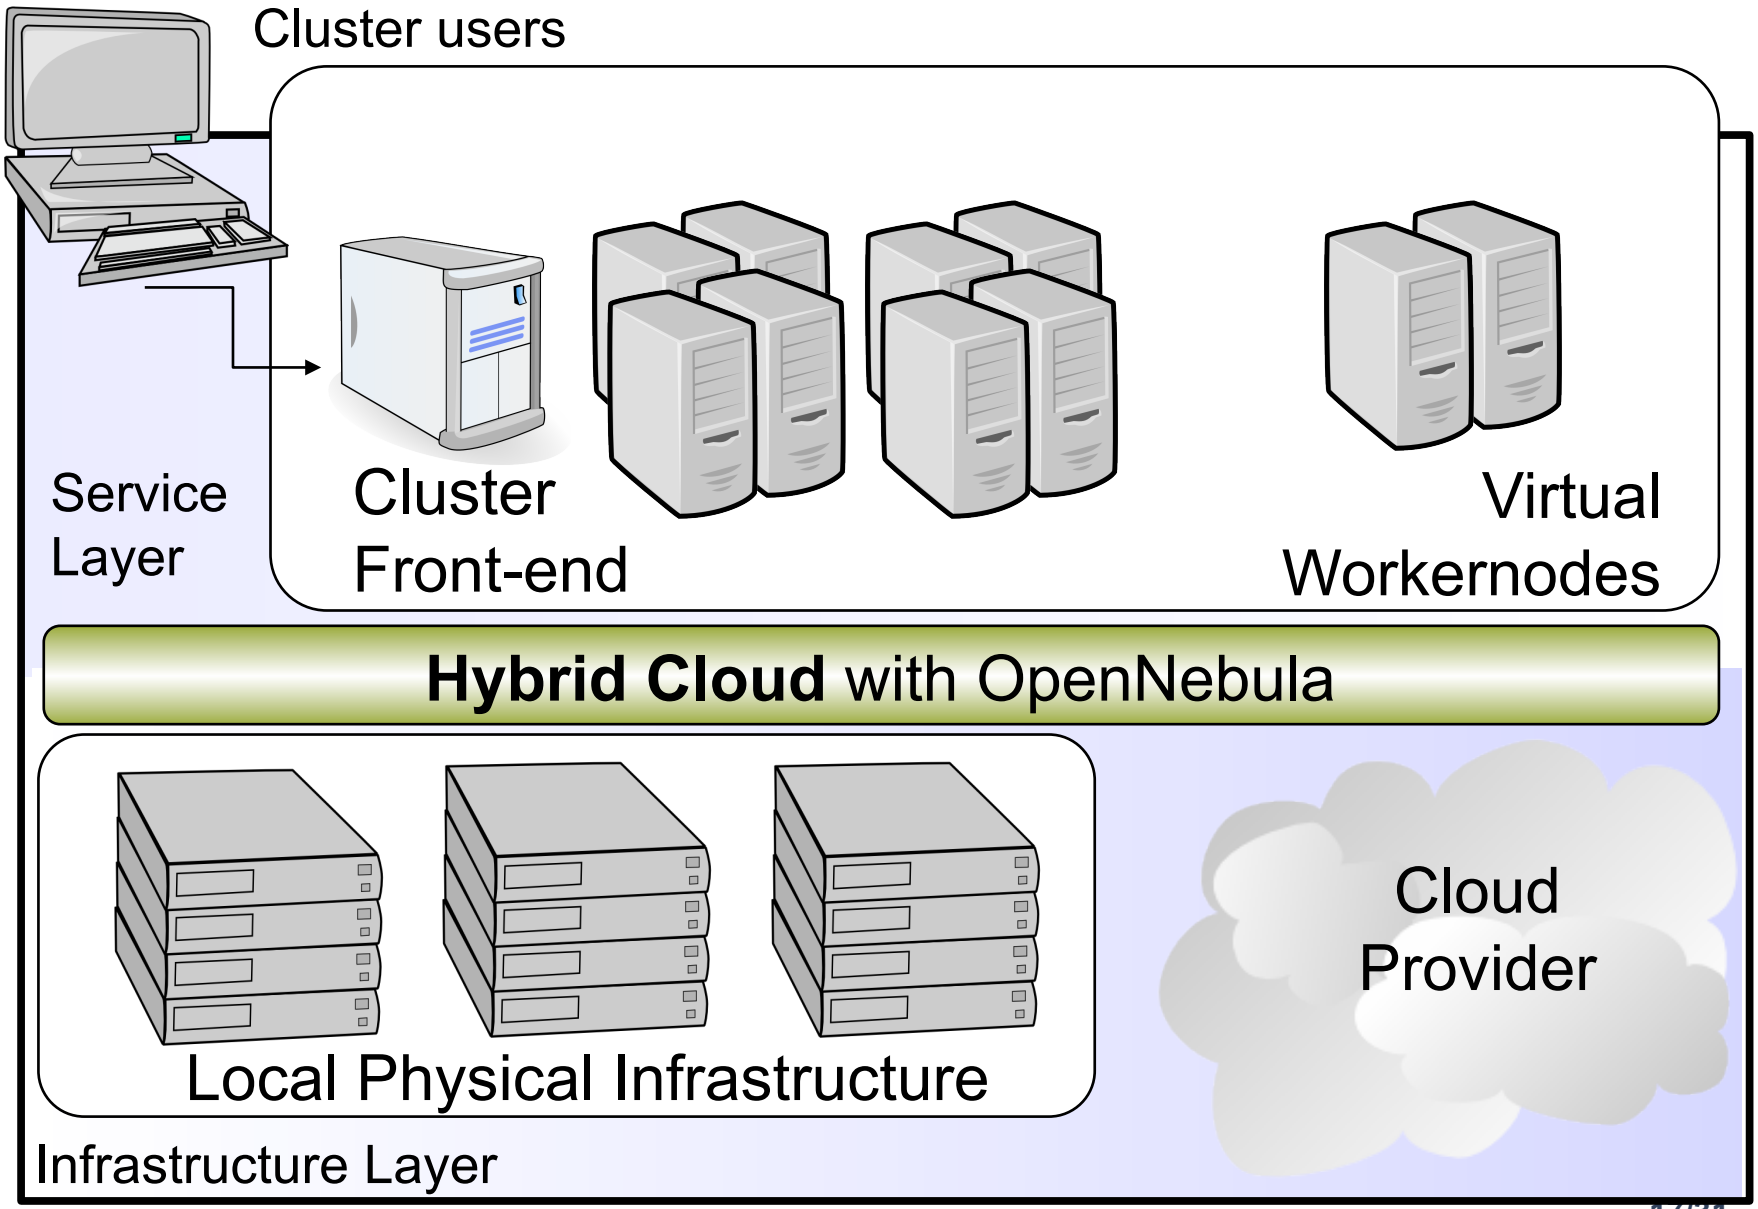

A Tool to Enhance Computing Infrastructures

Building Clouds with OpenNebula and its Application to Grid Computing

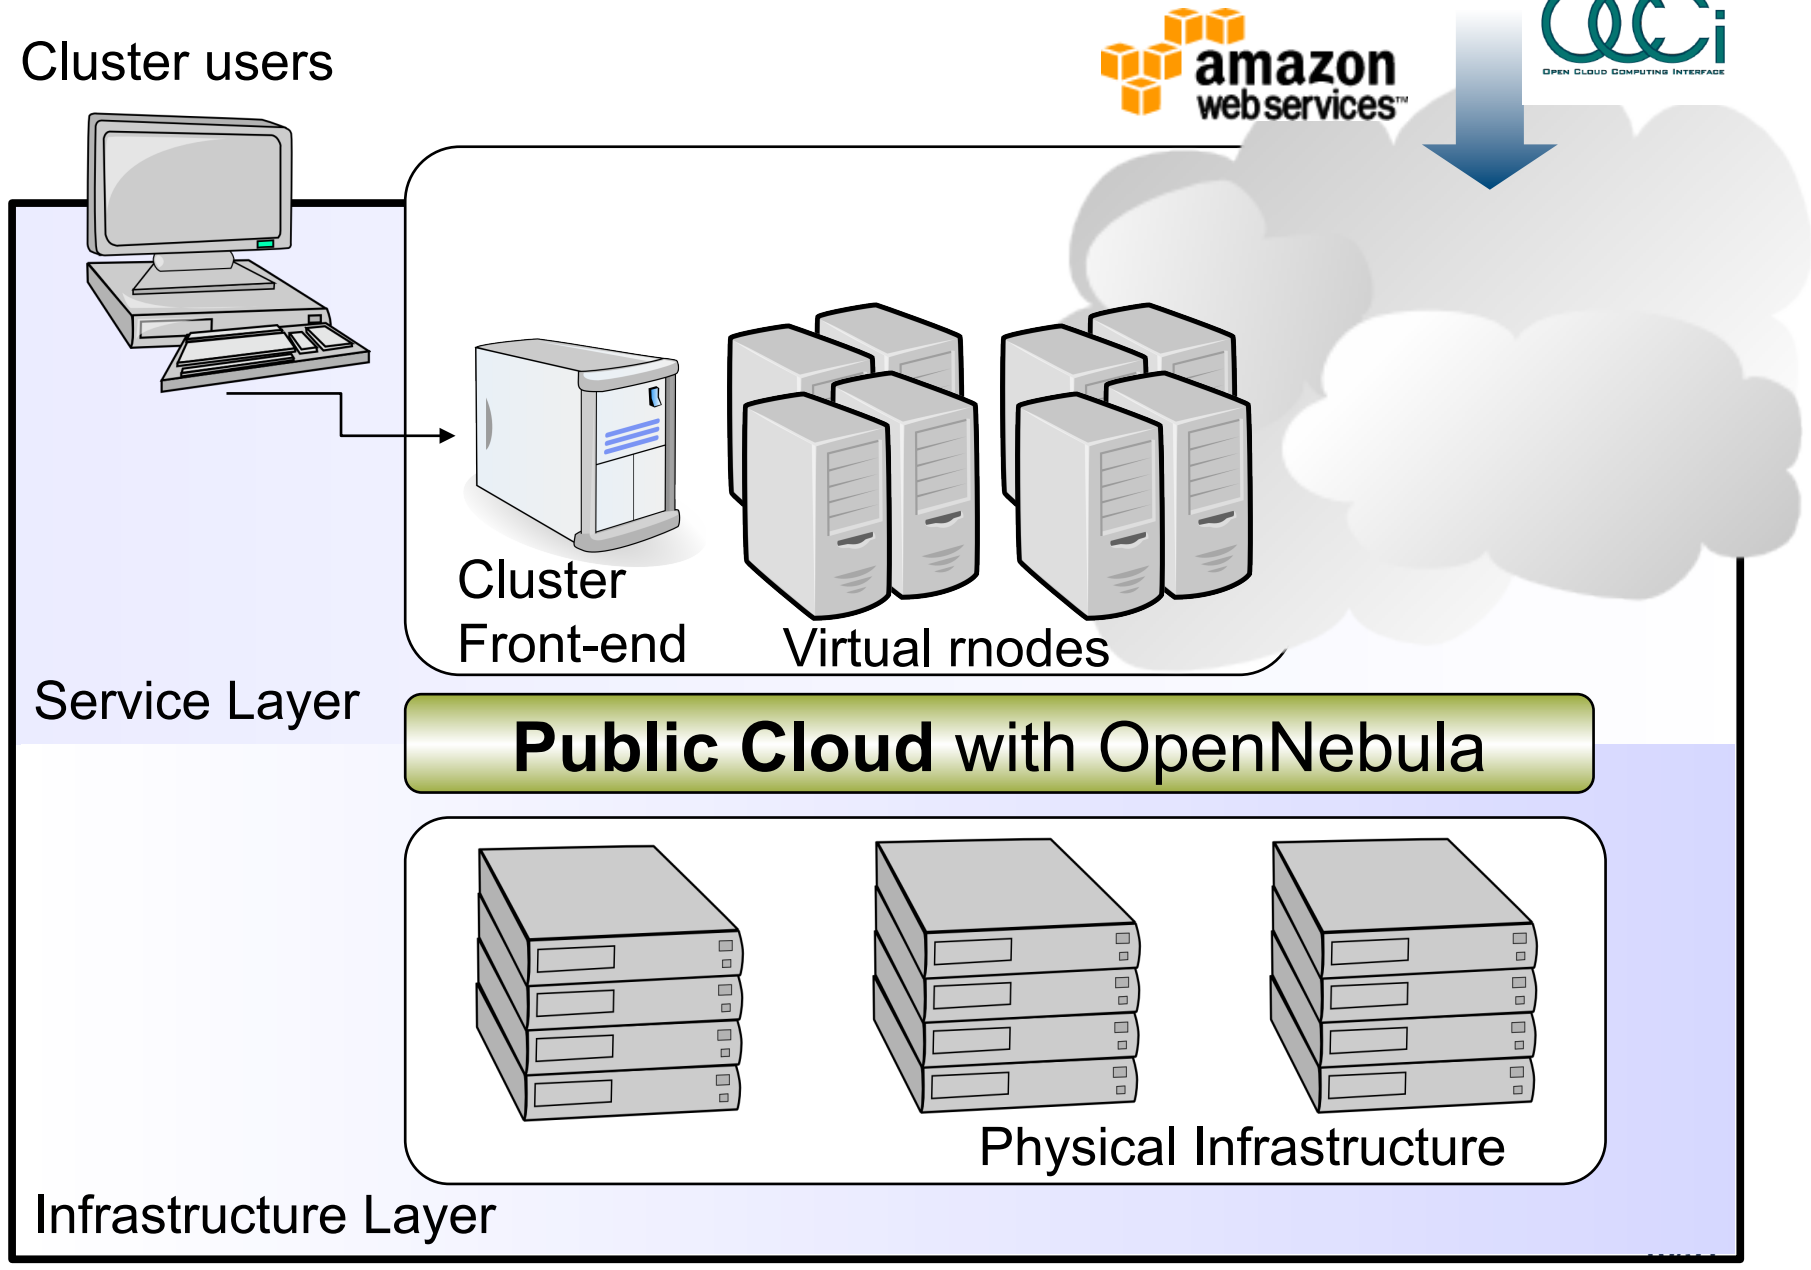

dsa-research.org

A Tool to Enhance Computing Infrastructures

Building Clouds with OpenNebula and its Application to Grid Computing

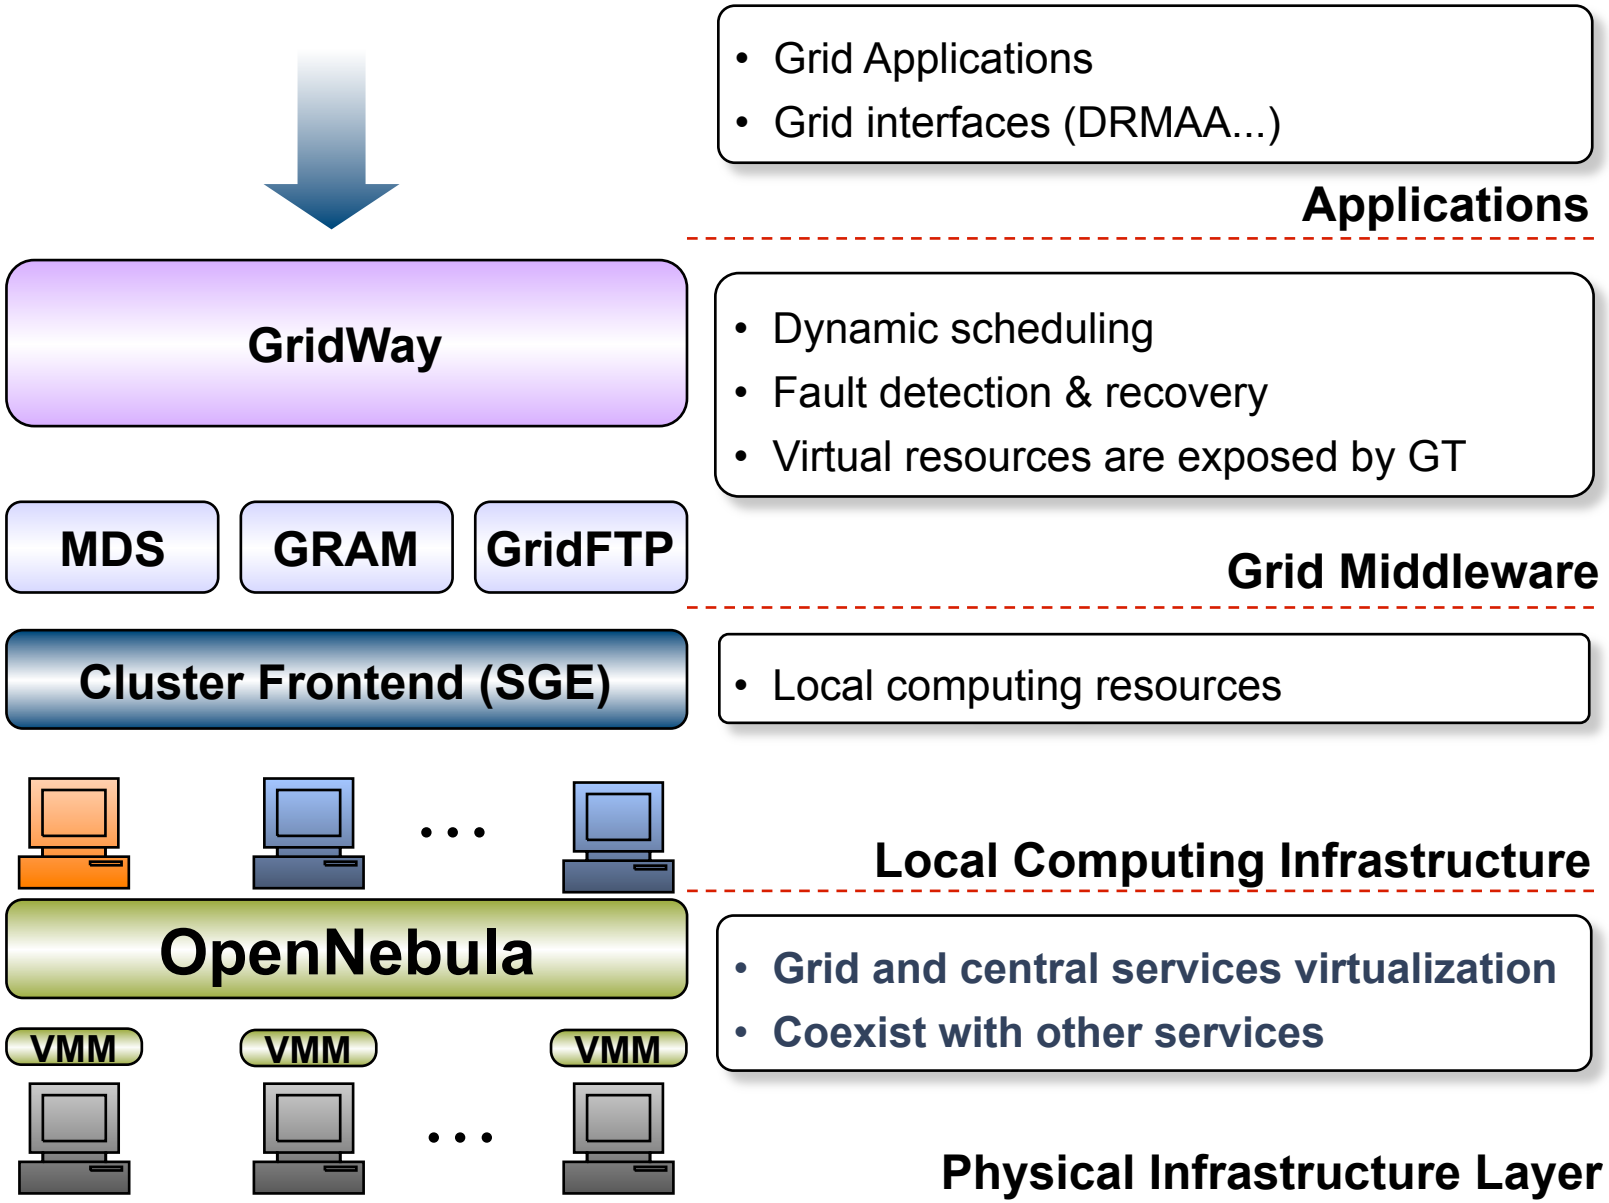

dsa-research.org

A Tool to Enhance Computing Infrastructures

Building Clouds with OpenNebula and its Application to Grid Computing

Benefits of Cloud for Existing Grid Infrastructures

- Easy support for VO-specific worker nodes
- Reduce *gridification* cycles
- Dynamic balance of resources between VO's and so maximize utility
- Fault tolerance of key infrastructure components
- Easier deployment and testing of new middleware distributions
- Distribution of pre-configured components
- Cheaper development nodes
- Simplified training machines deployment
- Simplified operation of grid sites
- Performance partitioning between local and grid services

Solve many of the obstacles to Grid adoption

A Tool to Enhance Computing Infrastructures

Building Clouds with OpenNebula and its Application to Grid Computing

Deployment Cases: Private Cloud to Support Grid Site

- The Dgrid Resource Center Ruhr (DGRZR) runs an OpenNebula private cloud on 248 blades and 1,984 cores with Xen
- OpenNebula is used to support the execution of a virtualized Grid site in D-Grid and EGEE

Deployment Cases: Public HPC Cloud

- SARA High Performance Computing Center uses OpenNebula in its new HPC Cloud service on 128 cores across 16 servers with KVM
- OpenNebula is used to support the execution of virtual clusters and HPC applications
- Authors of the OpenNebula Management Console

A Tool to Enhance Computing Infrastructures

Building Clouds with OpenNebula and its Application to Grid Computing

Deployment Cases: Private Cloud to Support Batch Farm

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

Up to 7,500 VMs on 400 hosts (3,200 cores) running Xen

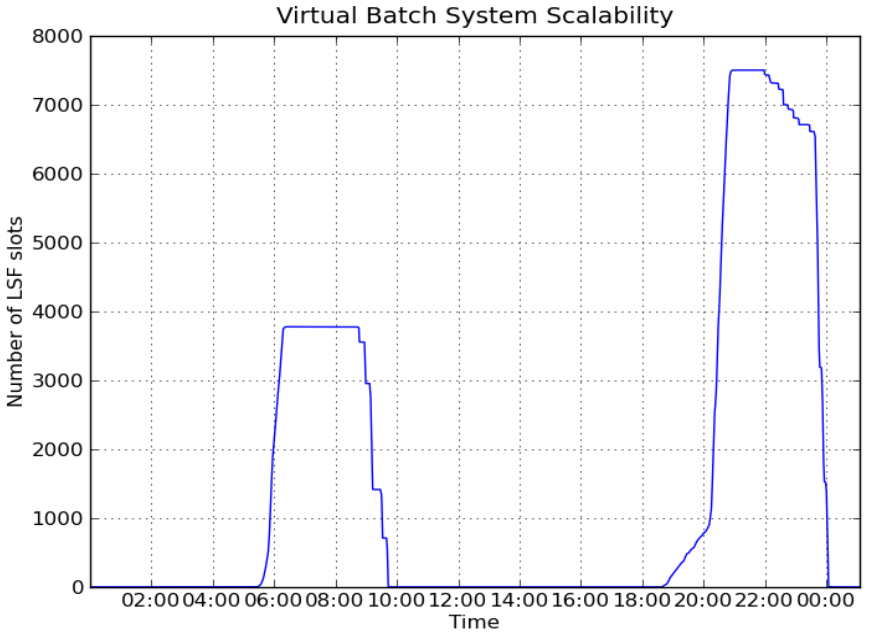

A Tool for Innovation

Building Clouds with OpenNebula and its Application to Grid Computing

European Projects on Cloud Computing Infrastructures

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

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

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

A Tool for Innovation: The Enabling Software Artefacts

Building Clouds with OpenNebula and its Application to Grid Computing

RESERVOIR **www.reservoir-fp7.eu**

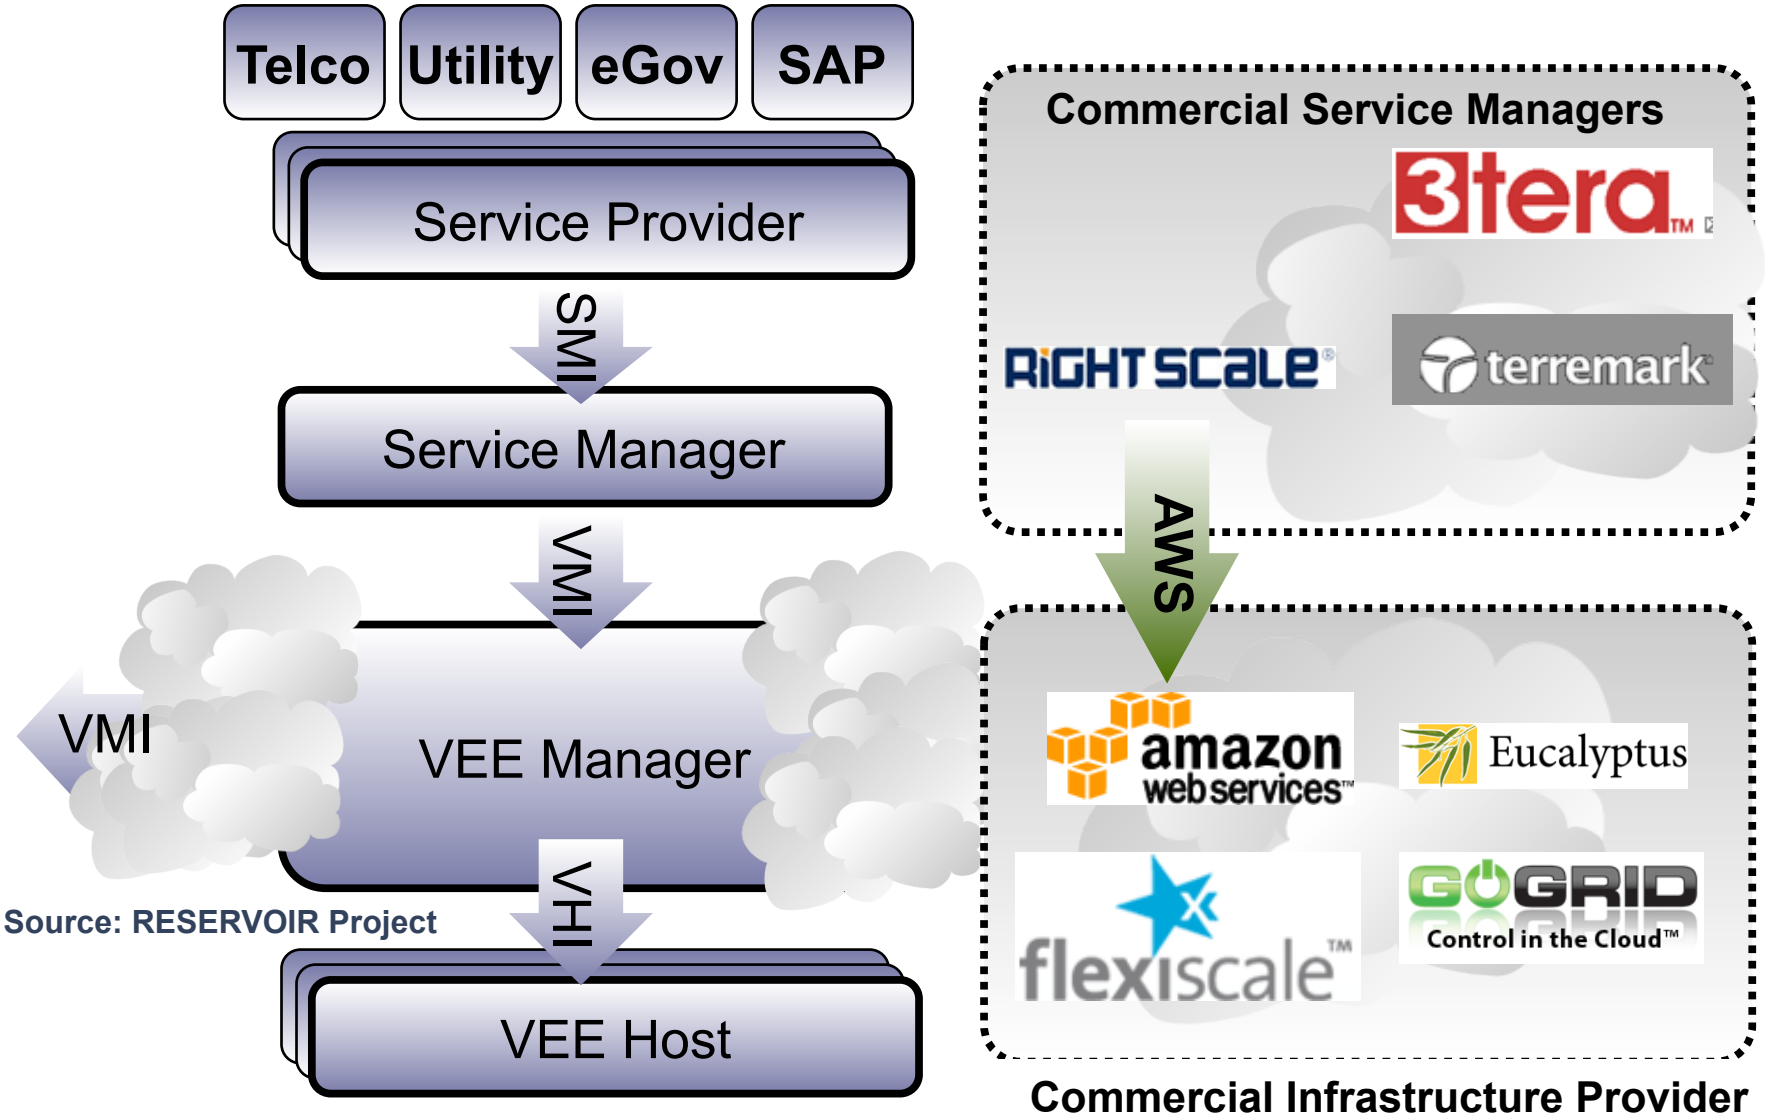

Source: RESERVOIR Project

A Tool for Innovation: Cloud for Service Experimentation

Building Clouds with OpenNebula and its Application to Grid Computing

Building Service Testbeds on FIRE

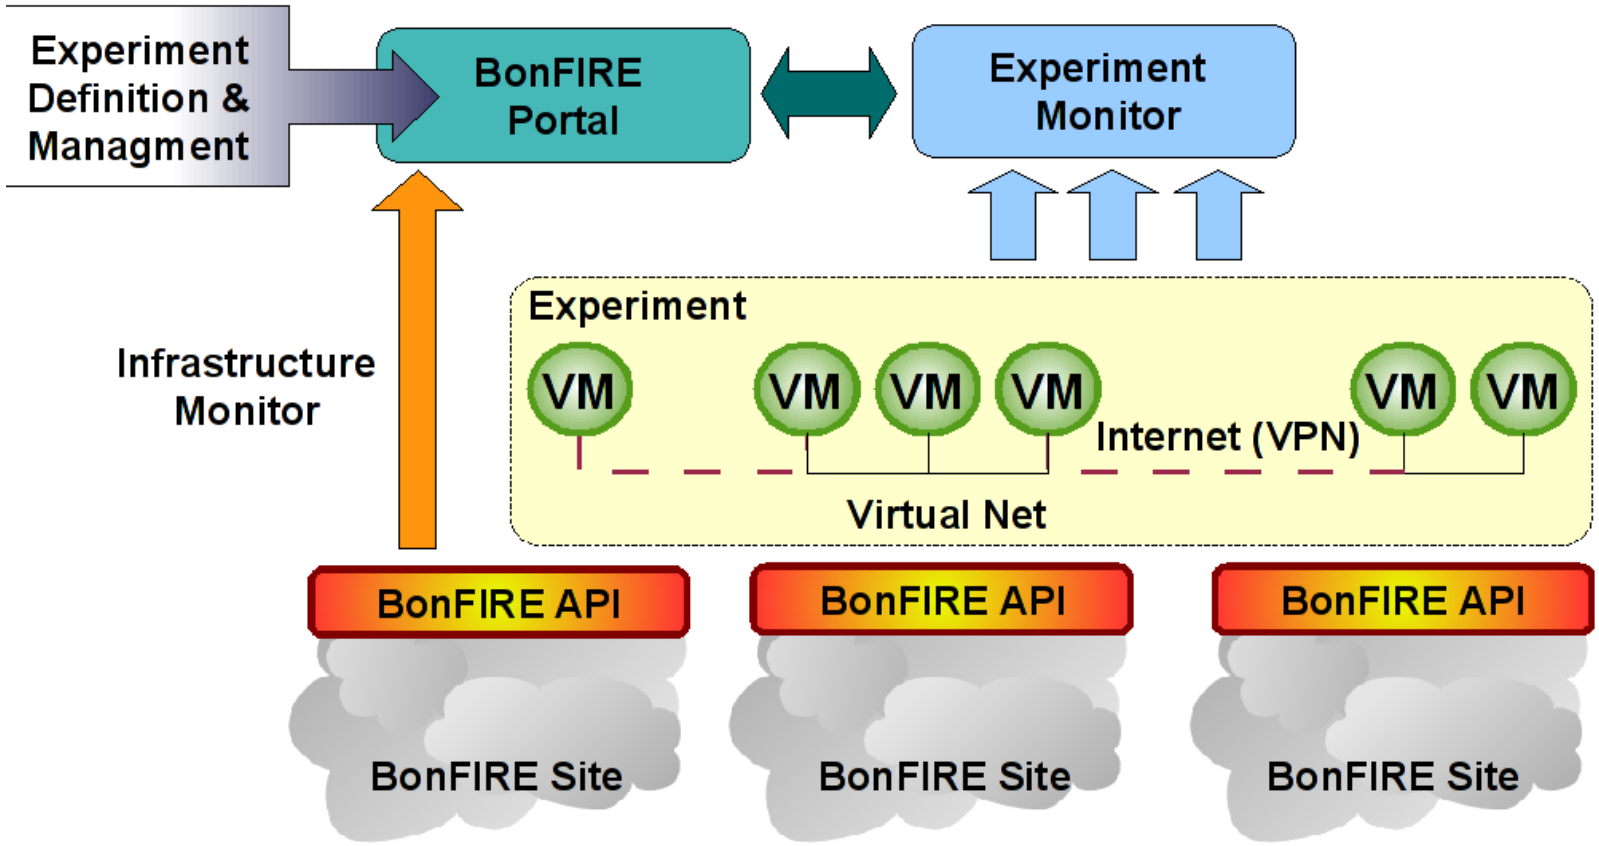

Source: BonFIRE Project

dsa-research.org

A Tool for Innovation: Enhancing Grid with Cloud

Building Clouds with OpenNebula and its Application to Grid Computing

StratusLab

Vision

- Grid and cloud embody **complementary computing models** that will coexist and cooperate in existing and future e-infrastructures

Aim

- Incorporate **cloud innovation into existing Grid infrastructures** to:
  - **Simplify** and **optimize** its use and operation, providing a more **flexible**, **dynamic** computing environment for scientists.
  - **Enhance existing infrastructures with "IaaS" cloud paradigms**

Evolutionary Approach

- **Complement existing services**, being fully transparent to upper layers
- **Existing Grid middleware would continue** to provide the glue to federate the distributed resources and the services for high-level job and data management
- **Address the emerging IaaS cloud-like usage patterns**

Service Centred Project driven to support production infrastructures

- Integration, distribution, testing and maintenance the **StratusLab Toolkit**
- **Management of images** (trust, provenance…)

A Tool for Innovation: Enhancing Grid with Cloud

Building Clouds with OpenNebula and its Application to Grid Computing

www.stratuslab.org

StratusLab

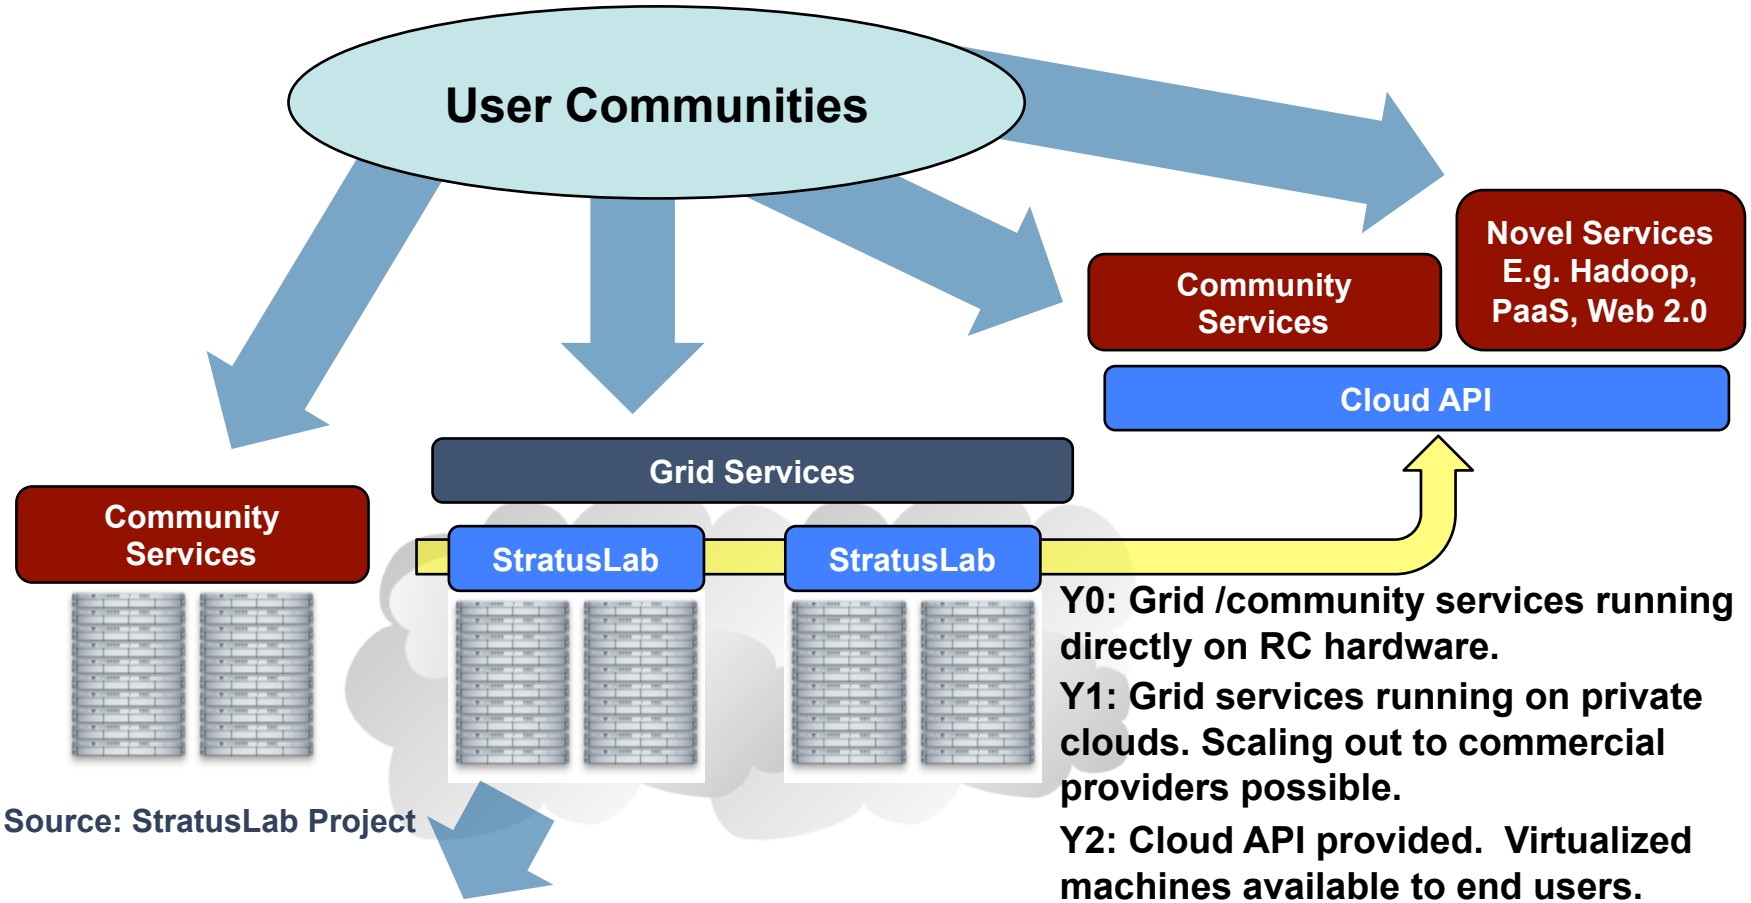

Y0: Grid /community services running directly on RC hardware.

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.

Outlook

About the Short-term Roadmap (2 months): v1.6

Feature	New Function
Scalability, Reliability and High Availability	• Support fro MySQL in the back-end • Unit-testing of the core • HTTP back-end
Functionality	• Image repository • Support for multiple clusters • CLI for accounting and billing support
Cloud Interfaces	• Improve compatibility with EC2 ecosystem

About the Medium-term Roadmap

- **Projects** funding OpenNebula
- **Community**

Funding

- New European Projects ensure the development and maintenance of OpenNebula until end of 2013
- C12G Labs also contributes to the sustainability of the open-source community

Long-term Sustainability and Commercial Support

Building Clouds with OpenNebula and its Application to Grid Computing

C12G.com

dsa-research.org

Thanks

Funding Agencies

- **European Commission**: RESERVOIR 2008-2011, EU agreement 215605
- **Ministry Science&Innovation**: HPCcloud 2010-2012, MICINN TIN2009-07146
- **Community of Madrid**: MEADIANET 2010-2013 CAM S2009/TIC-1468

Other Sponsors

- **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!

More Information

More info, downloads, mailing lists at

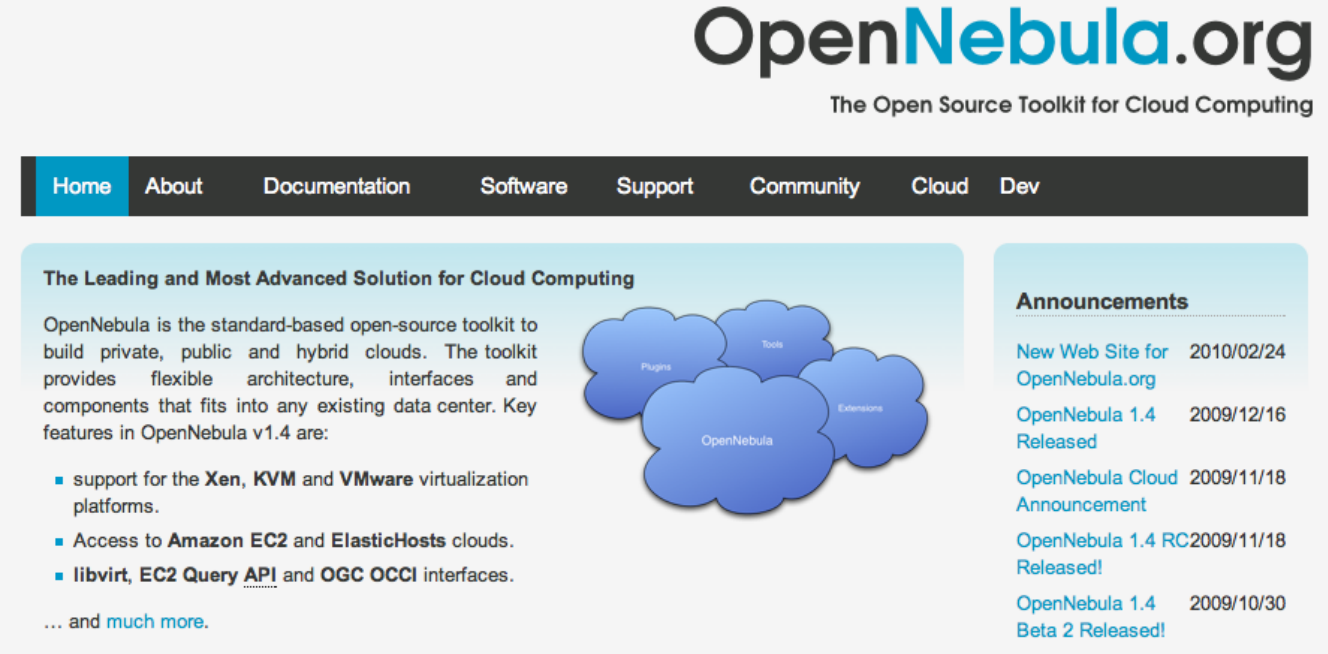

Research References

- 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)