

Euro-Par 2010
1st September 2010
Ischia, Naples, Italy

Innovation in Cloud Computing Architectures

Ignacio M. Llorente

dsa-research.org

Distributed Systems Architecture Research Group
Universidad Complutense de Madrid



Types of Cloud Services

Innovation in Cloud Computing Architectures

Software as a Service

What

Who

On-demand access to any application

End-user (does not care about hw or sw)



Platform as a Service

Platform for building and delivering web applications

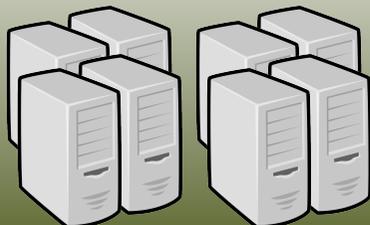
Developer (no managing of the underlying hw & sw layers)



Infrastructure as a Service

Raw computer infrastructure

System Administrator (complete management of the computer infrastructure)



Physical Infrastructure

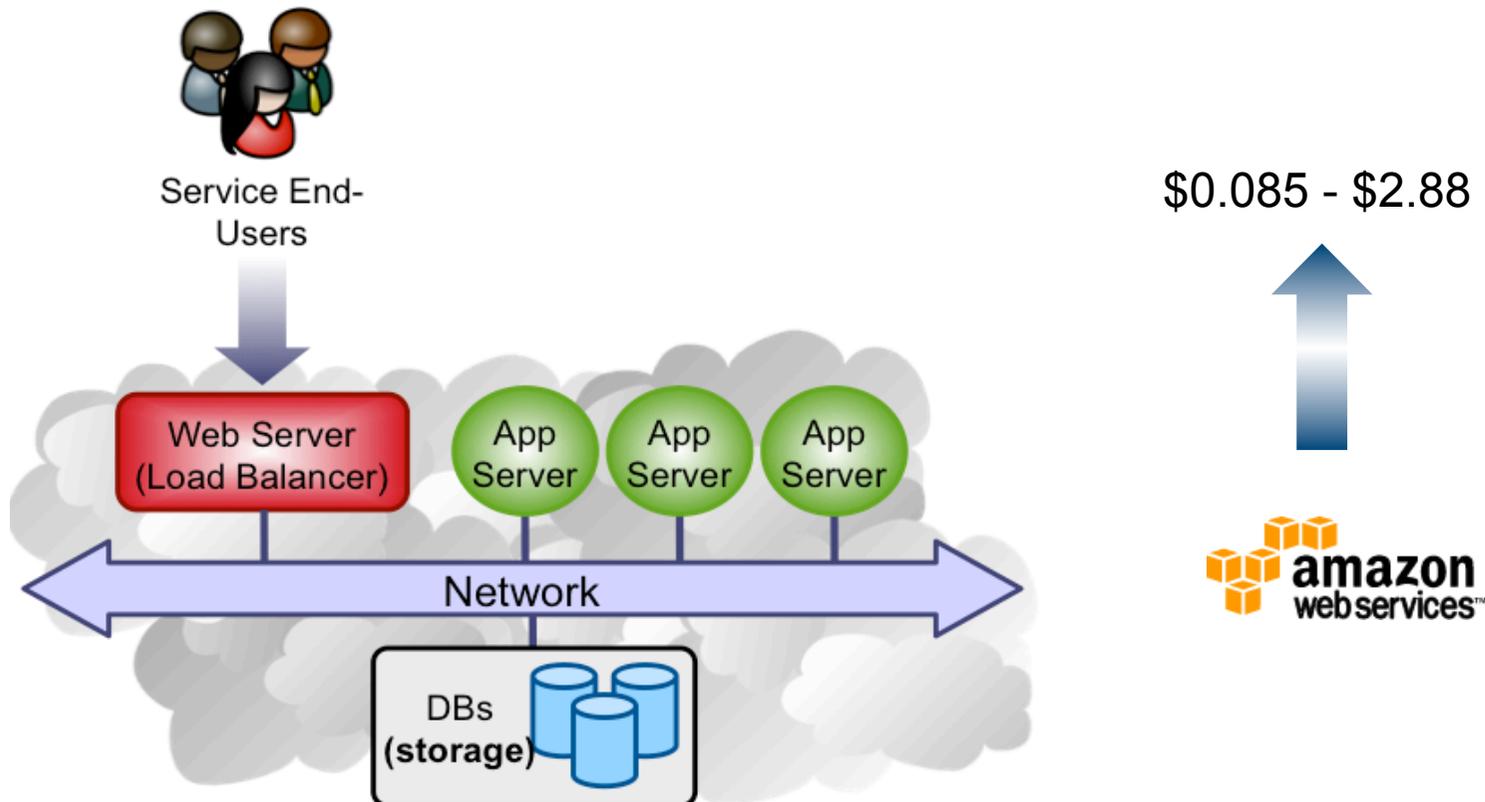


IaaS Cloud Computing

Innovation in Cloud Computing Architectures

Commercial Cloud Provider

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

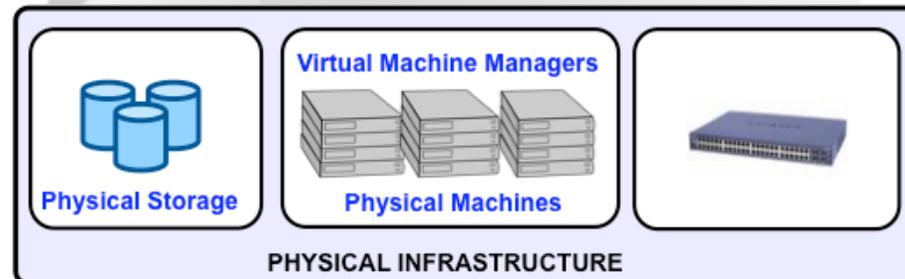


Transform your Infrastructure into a Cloud

Innovation in Cloud Computing Architectures

Building your Own Cloud

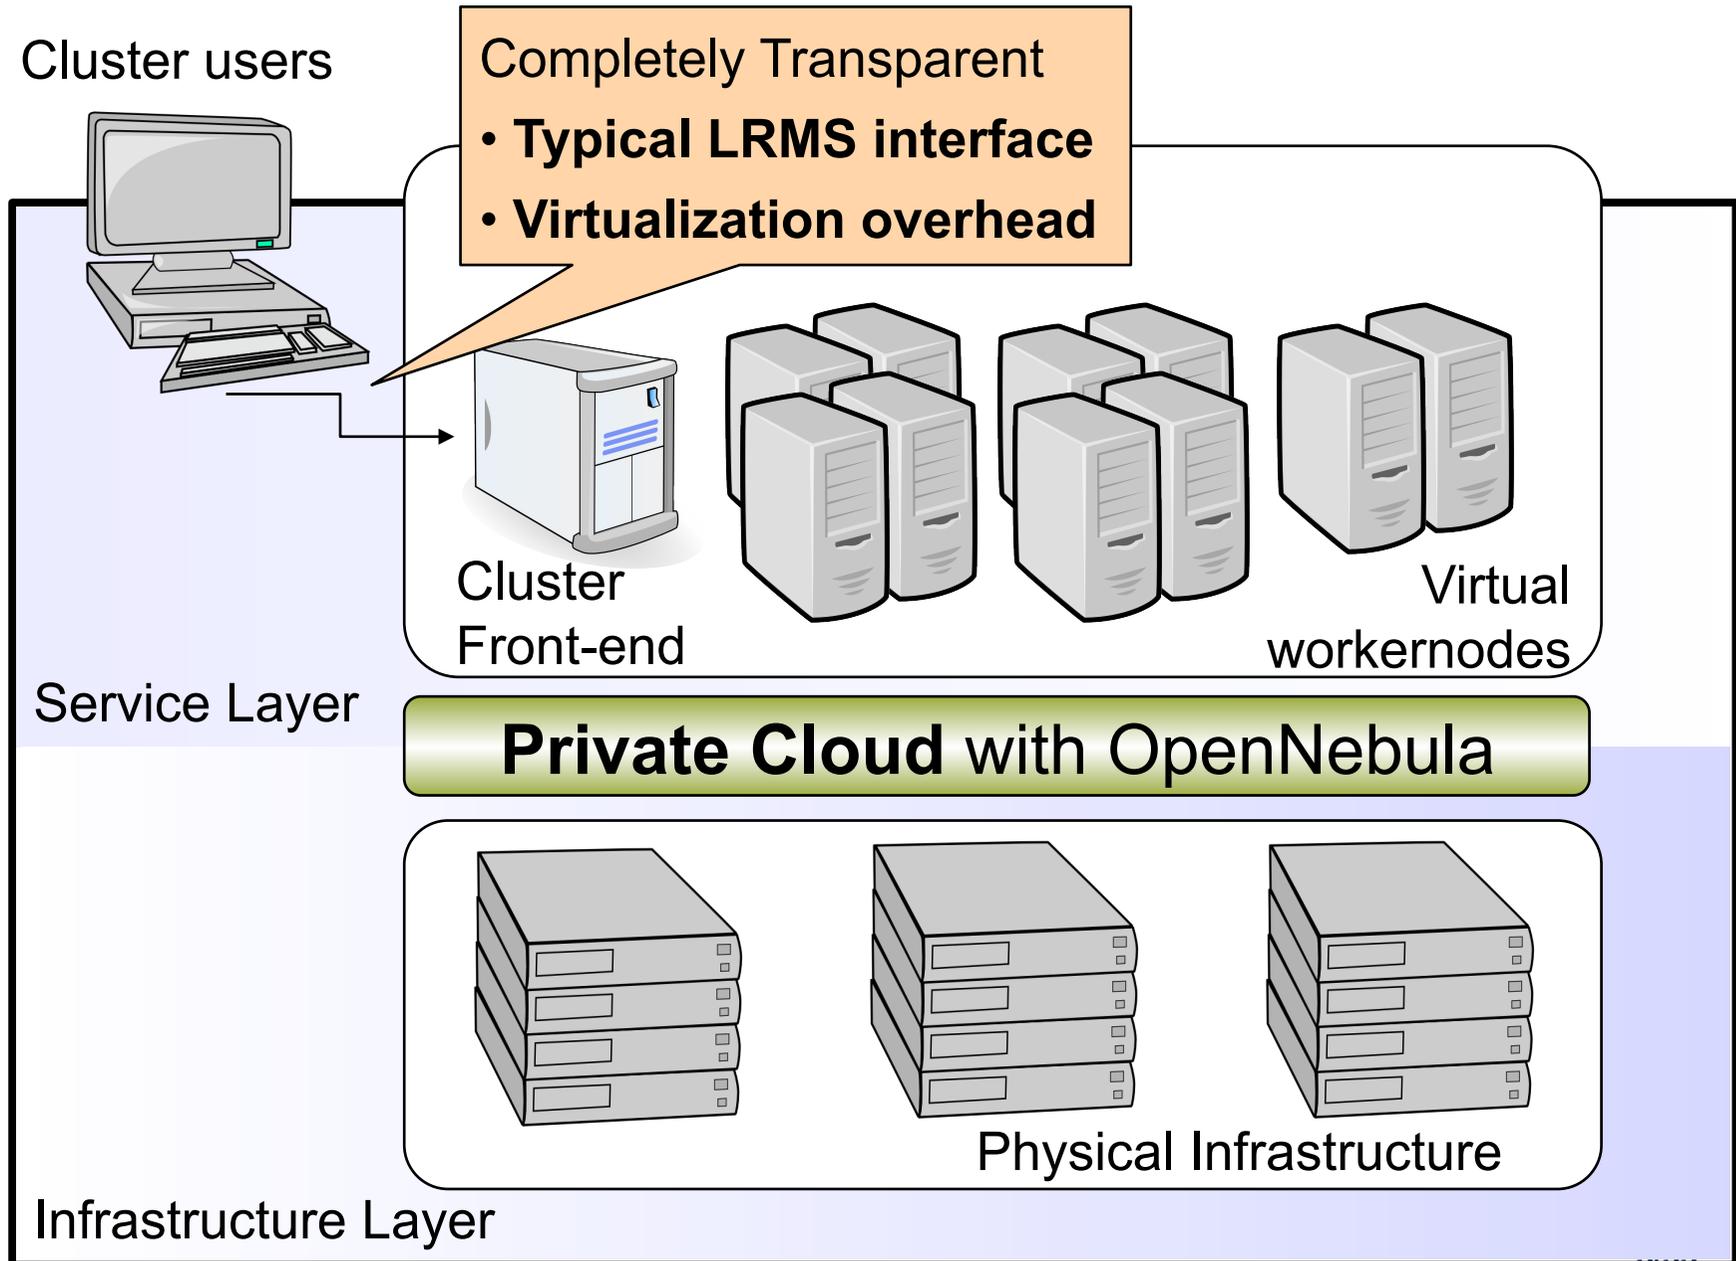
- **Optimize and Simplify Internal Operations**
 - **Centralized management** of all servers and services with dynamic resizing of infrastructure and dynamic allocation of capacity
 - **Higher utilization** and **operational saving** of existing resources with server consolidation and removal of application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





What is in for me?: A Cloud for Computing

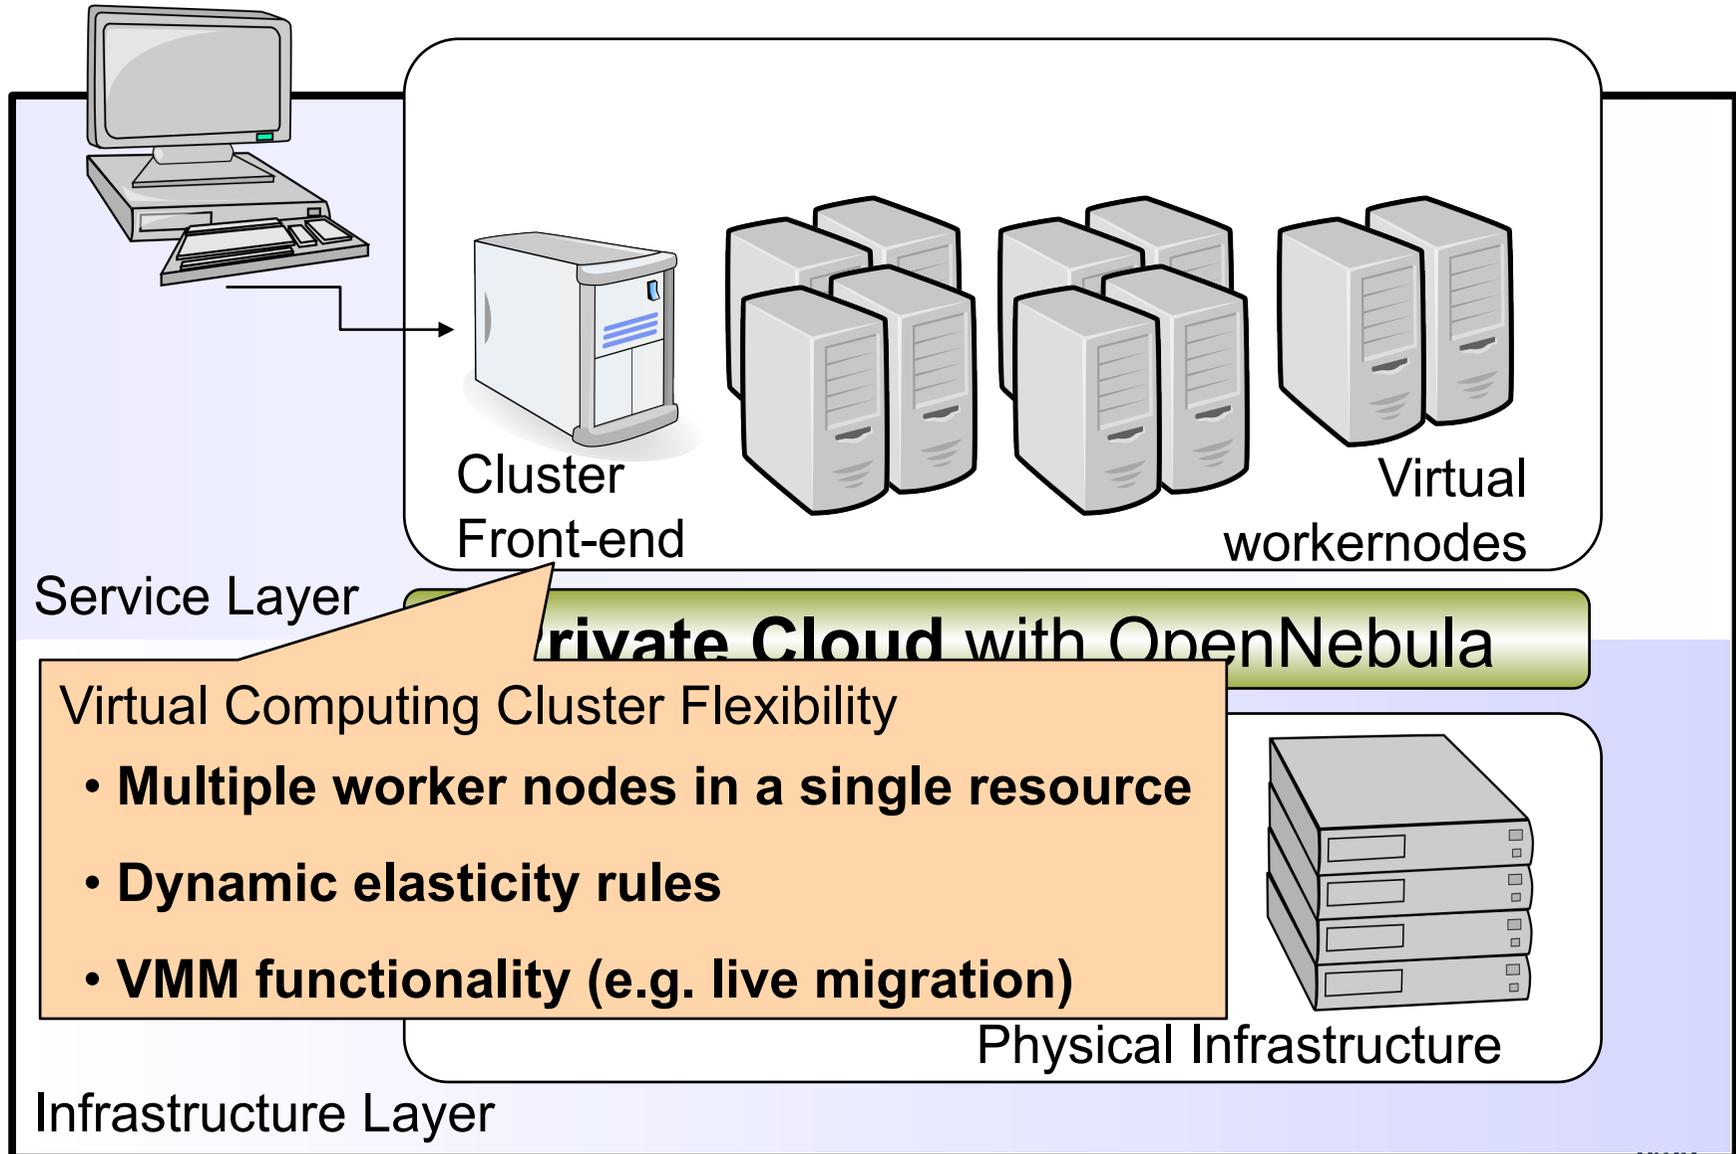
Innovation in Cloud Computing Architectures



What is in for me?: A Cloud for Computing

Innovation in Cloud Computing Architectures

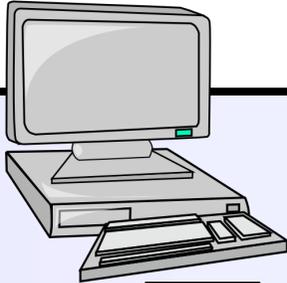
Cluster users



What is in for me?: A Cloud for Computing

Innovation in Cloud Computing Architectures

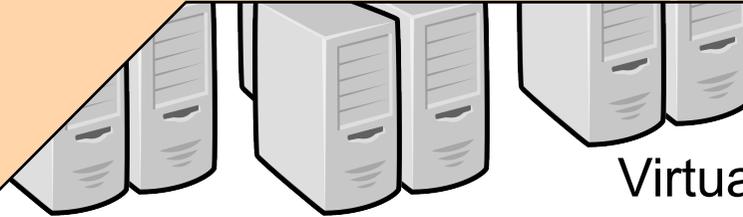
Cluster users



Physical Cluster Partitioning

- Performance partitioning (dedicated nodes)
- Isolate cluster workload
- Different physical clusters for different profiles

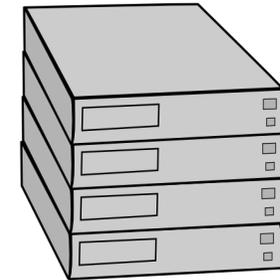
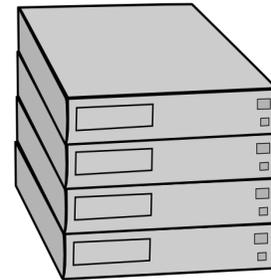
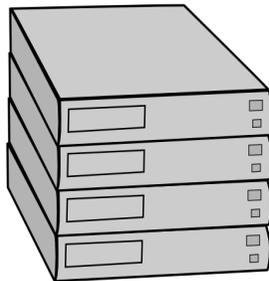
Cluster
Front



Virtual
workernodes

Service Layer

Private Cloud with OpenNebula



Physical Infrastructure

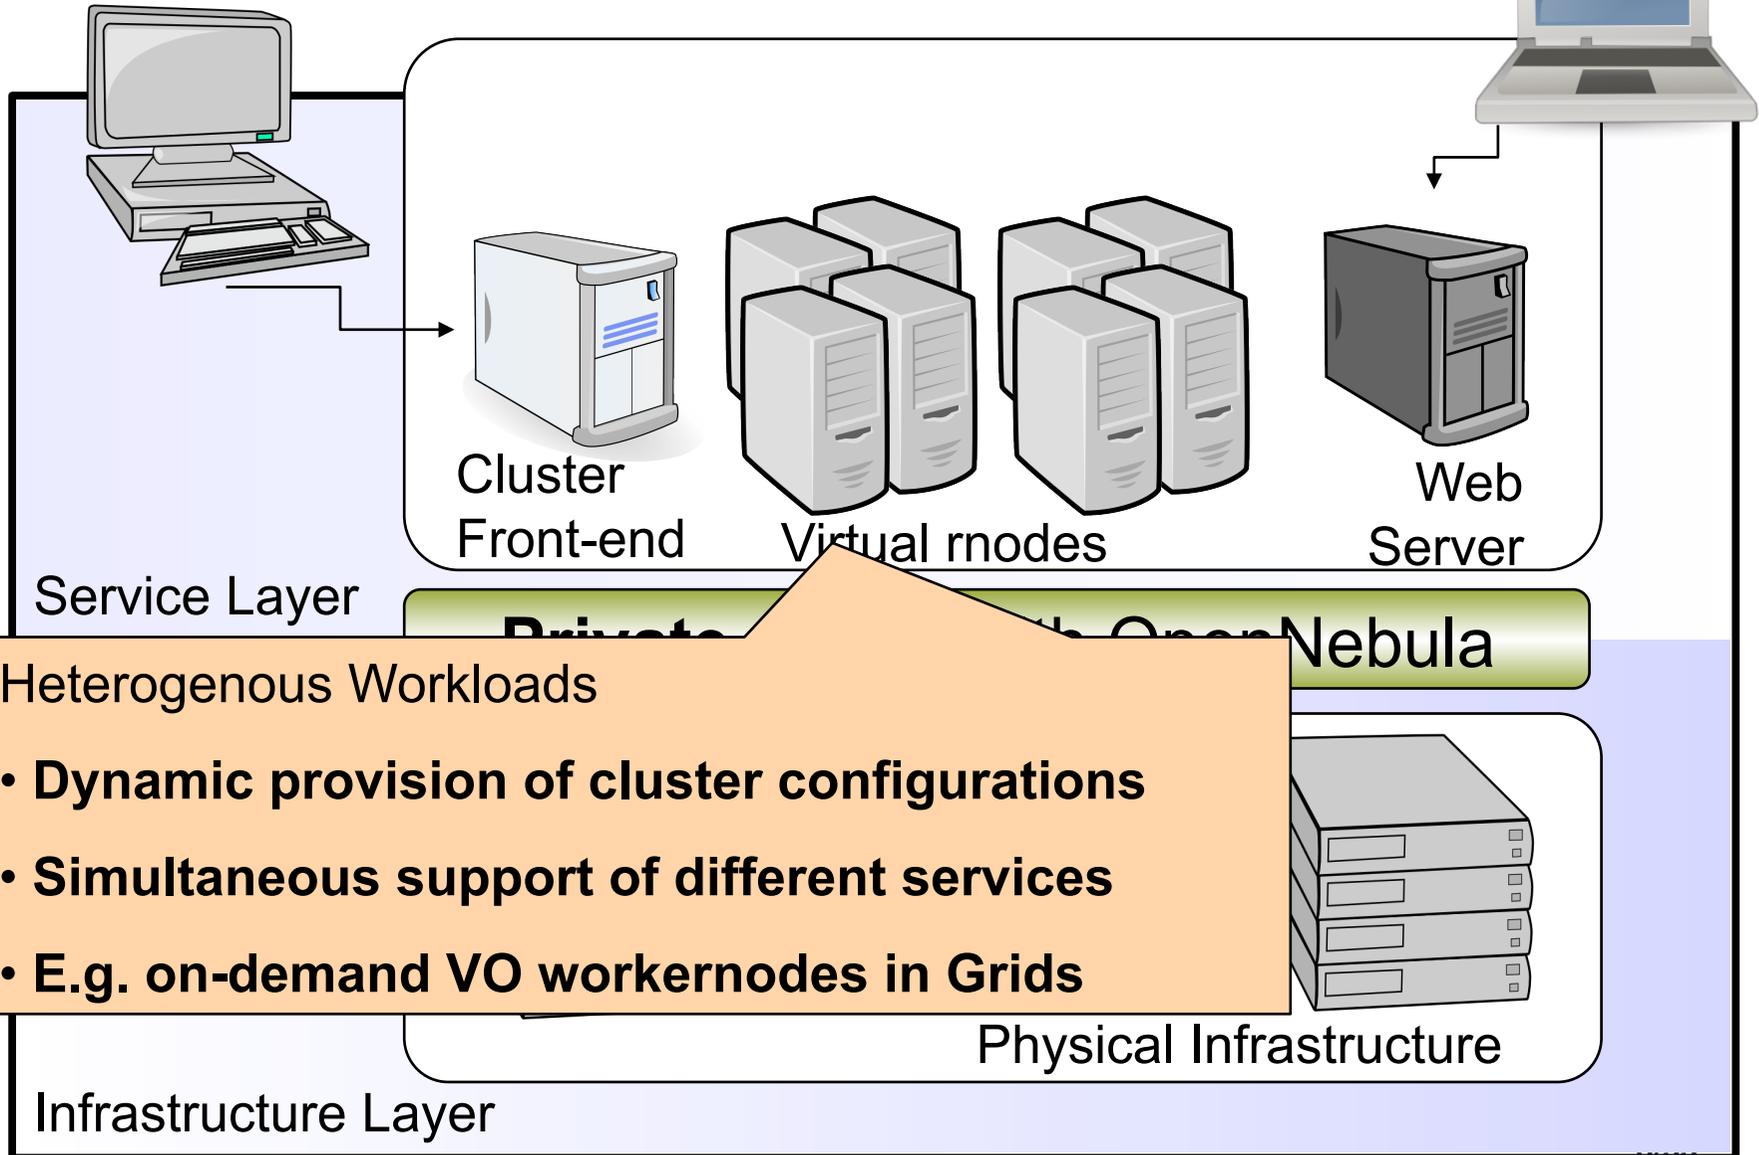
Infrastructure Layer

What is in for me?: A Cloud for Computing

Innovation in Cloud Computing Architectures

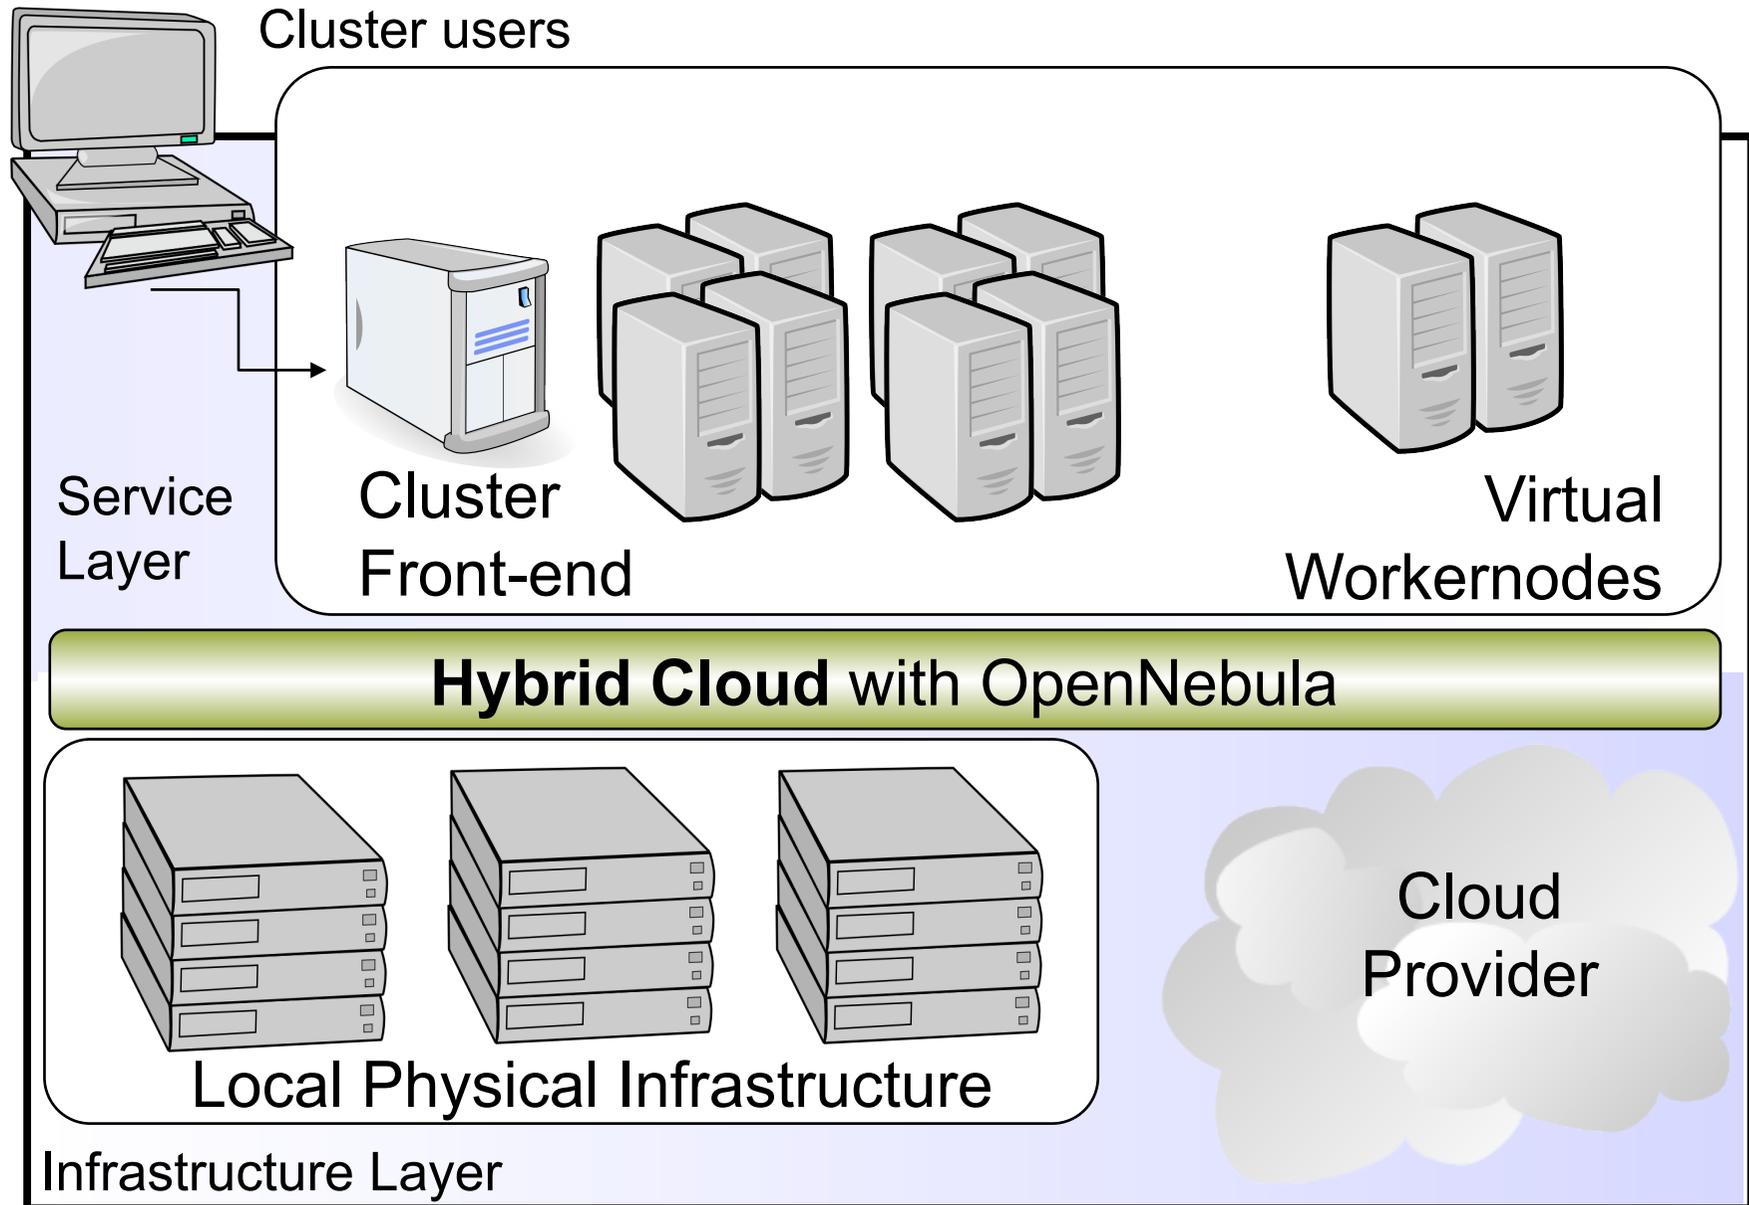
Cluster users

HTTP clients



What is in for me?: A Cloud for Computing

Innovation in Cloud Computing Architectures



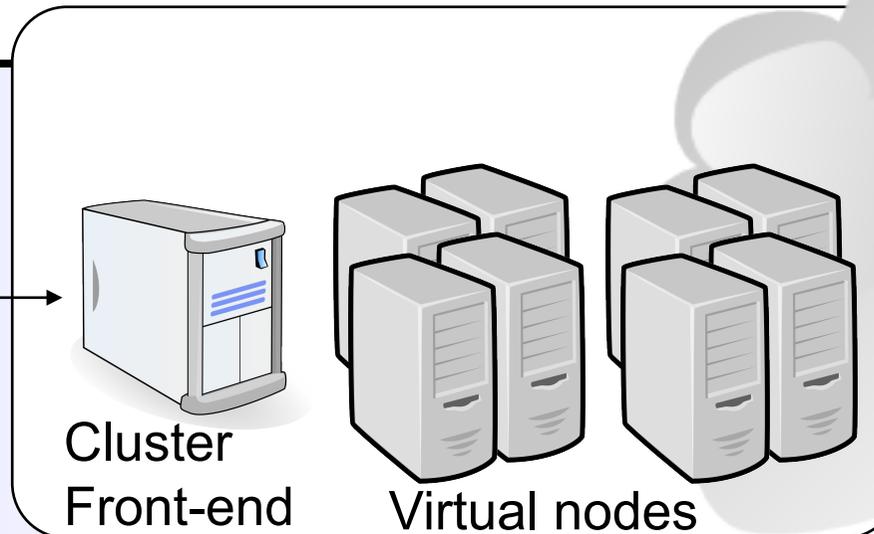
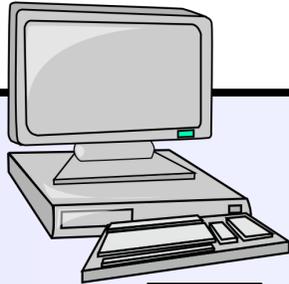


What is in for me?: A Cloud for Computing

Innovation in Cloud Computing Architectures

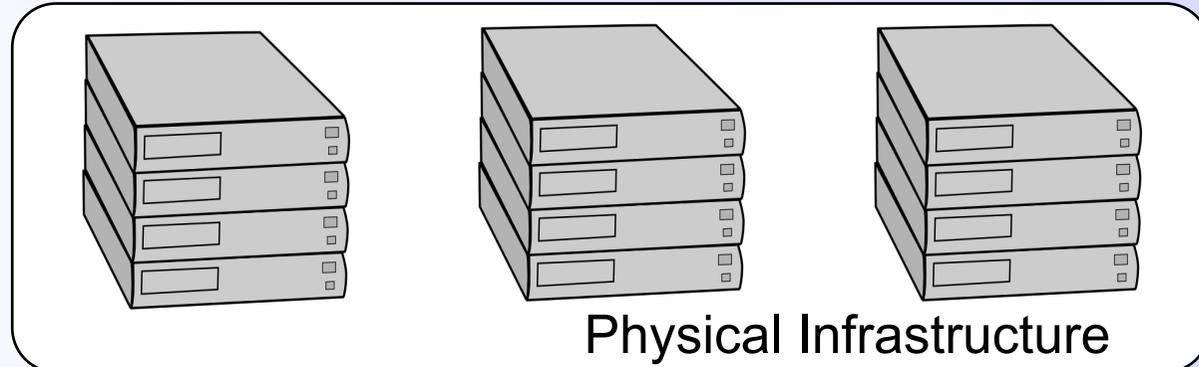


Cluster users



Service Layer

Public Cloud with OpenNebula



Infrastructure Layer

dsa-research.org

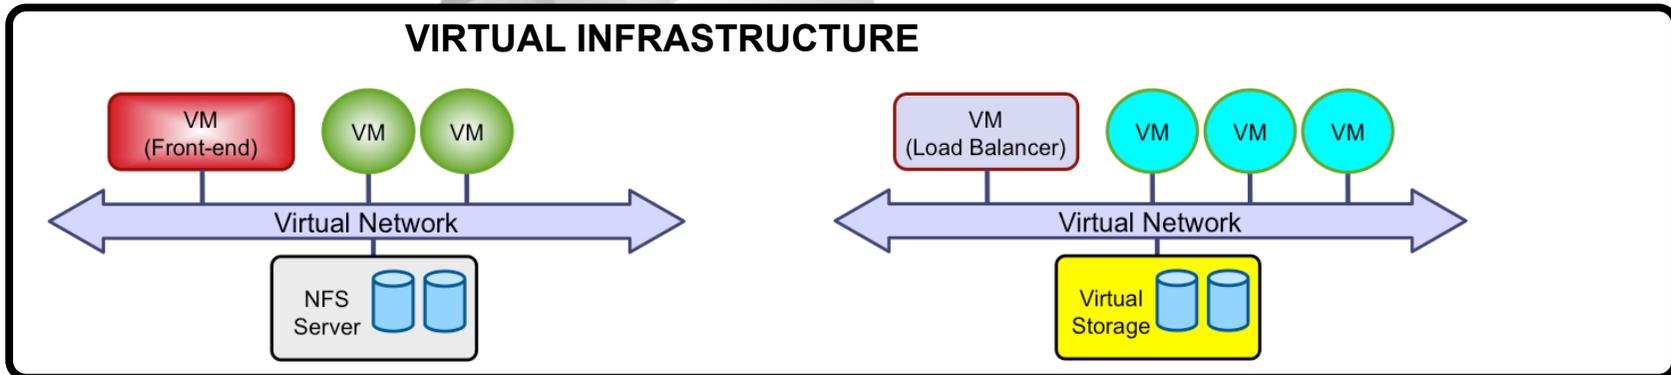


The Data Center is the Computer

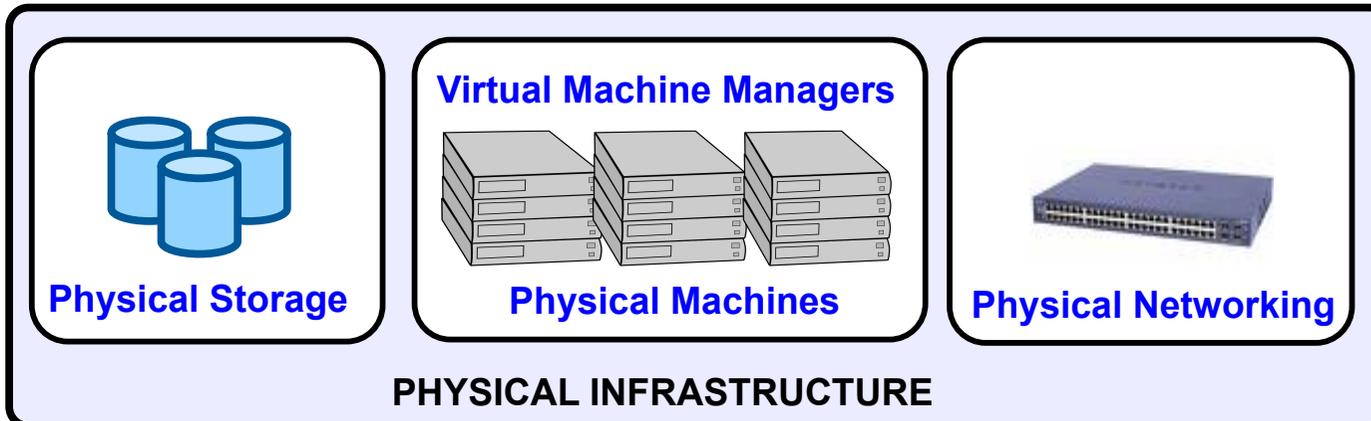
Innovation in Cloud Computing Architectures

Cloud Manager to Orchestrate the Complexity of a Datacenter

Service End-Users



OpenNebula.org





Contents

Innovation in Cloud Computing Architectures

dsa-research.org

Innovation in Cloud Computing Management

Addressing the technology challenges in cloud computing management from business use cases



Building a Cloud Infrastructure

OpenNebula as Cloud Enabling Technology



A Tool for Innovation and Research

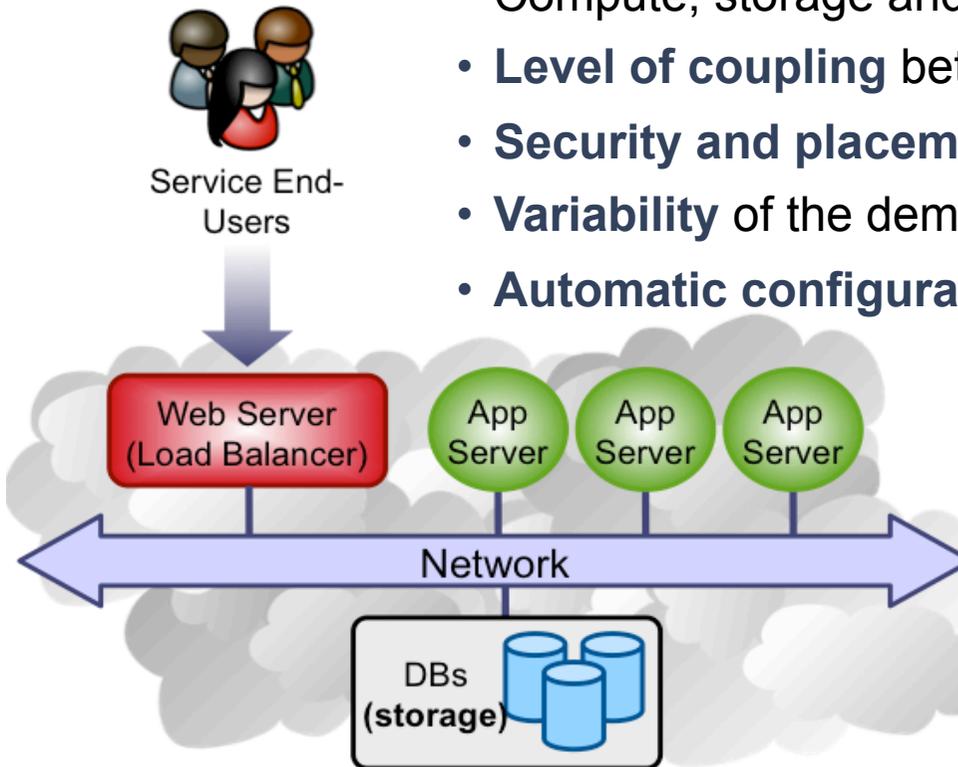
European Projects on Cloud Computing Infrastructures:
RESERVOIR, 4CaaS, StratusLab and BonFIRE

Innovations: The User Perspective

Innovation in Cloud Computing Architectures

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**
- **Variability** of the demand and **elasticity** of the services
- **Automatic configuration** of service instances

Service as Groups of VMs

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

Cloud Interface

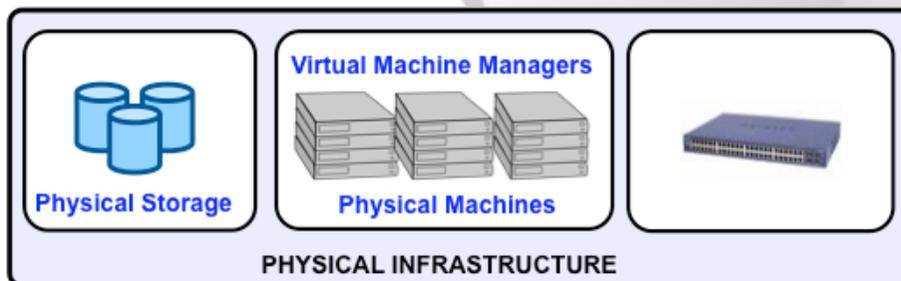
- Management of **VM images, VM instances, and virtual networks**
- **Support** for standard and common interfaces (OGF OCCl, Amazon EC2 and VMware vCloud)

Innovations: The Manager Perspective

Innovation in Cloud Computing Architectures

Efficient and Scalable 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
- Support for the definition of workload and resource-aware **allocation policies** such as consolidation (energy efficiency), load balancing, affinity-aware, capacity reservation, live migration...
- **Highly scalable back-end**



Scalable back-end

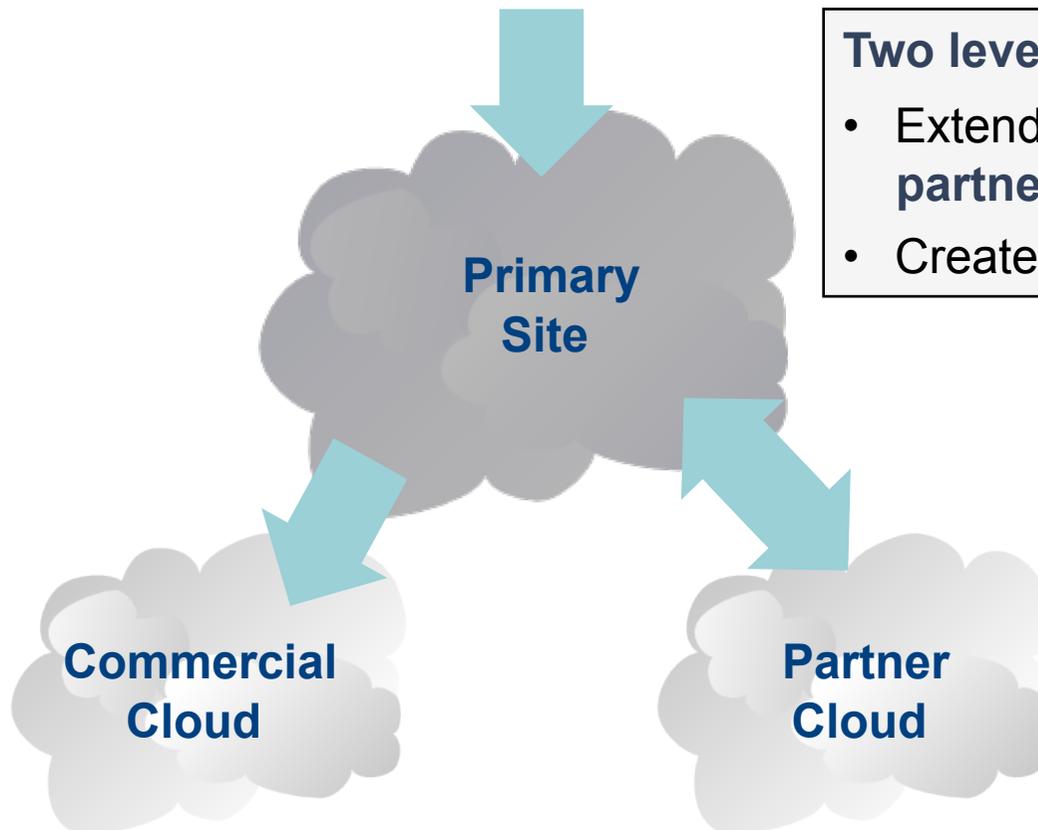
- Virtualization
- Storage
- Networking

Innovations: The Business Perspective

Innovation in Cloud Computing Architectures

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



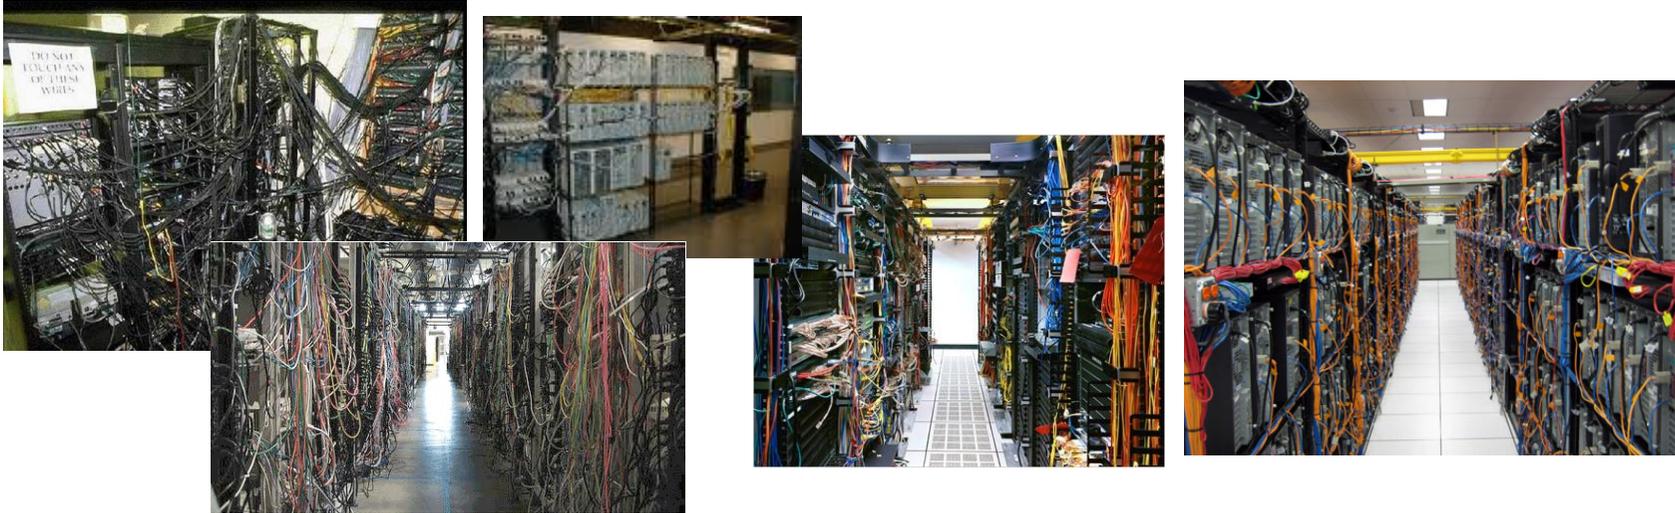
Two levels of Collaboration

- Extend the private cloud using both **partner and commercial clouds**
- Create a **federation of clouds**

Innovations: The Integrator Perspective

Innovation in Cloud Computing Architectures

From Heterogeneous and Ugly Data Centers...



... To Homogenous, Modular and Beautiful Data Center

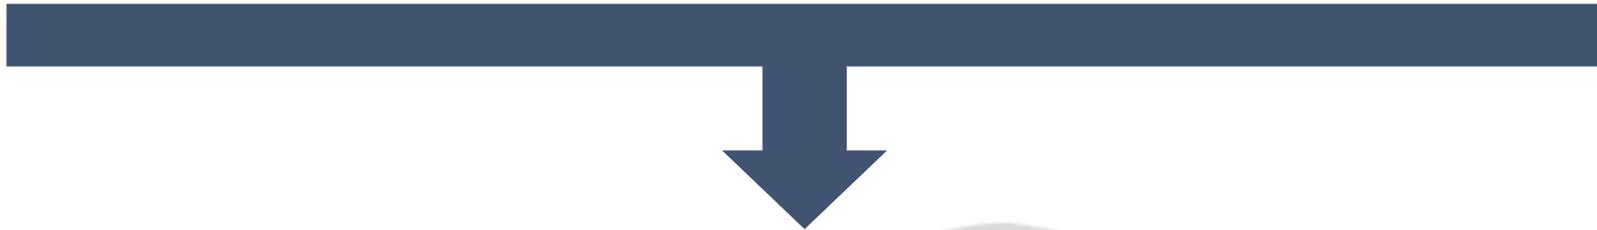


Innovations: The Integrator Perspective

Innovation in Cloud Computing Architectures

**Constraints from
Existing Infrastructure
and Processes**

**Requirements from
Usage and Deployment
Scenarios**



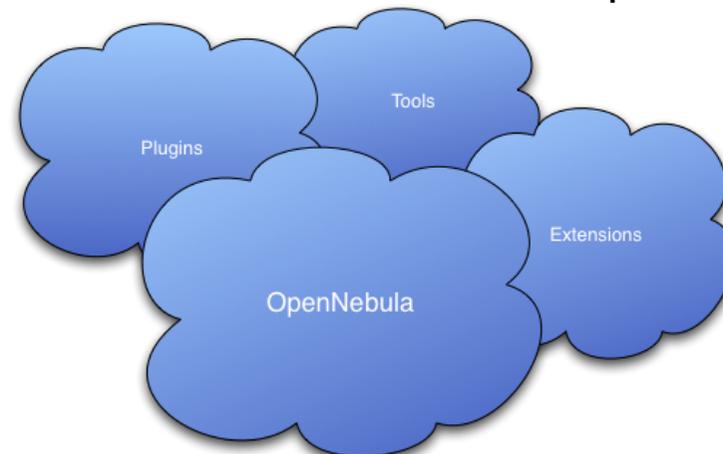
“One solution does not fit all requirements and constraints. There cannot be turnkey solutions for IaaS clouds”

Innovations: The Integrator Perspective

Innovation in Cloud Computing Architectures

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
- **Liberal open-source license**





Building a Cloud: Deployment Models

Innovation in Cloud Computing Architectures

Model	Definition	Examples of Deployment
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..)• Regional clouds to address regulatory or latency issues
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

Building a Cloud: The OpenNebula Toolkit

Innovation in Cloud Computing Architectures

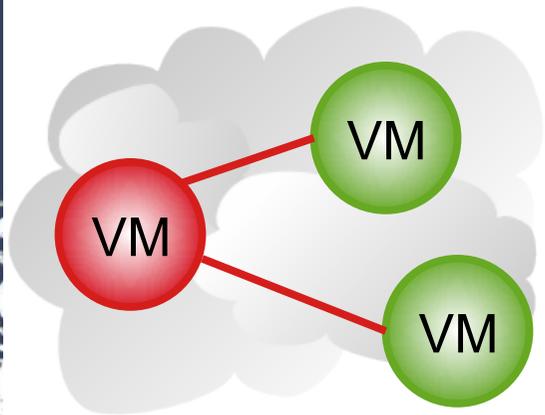
Flagship International Projects in Cloud Computing

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



Open-source Framework

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



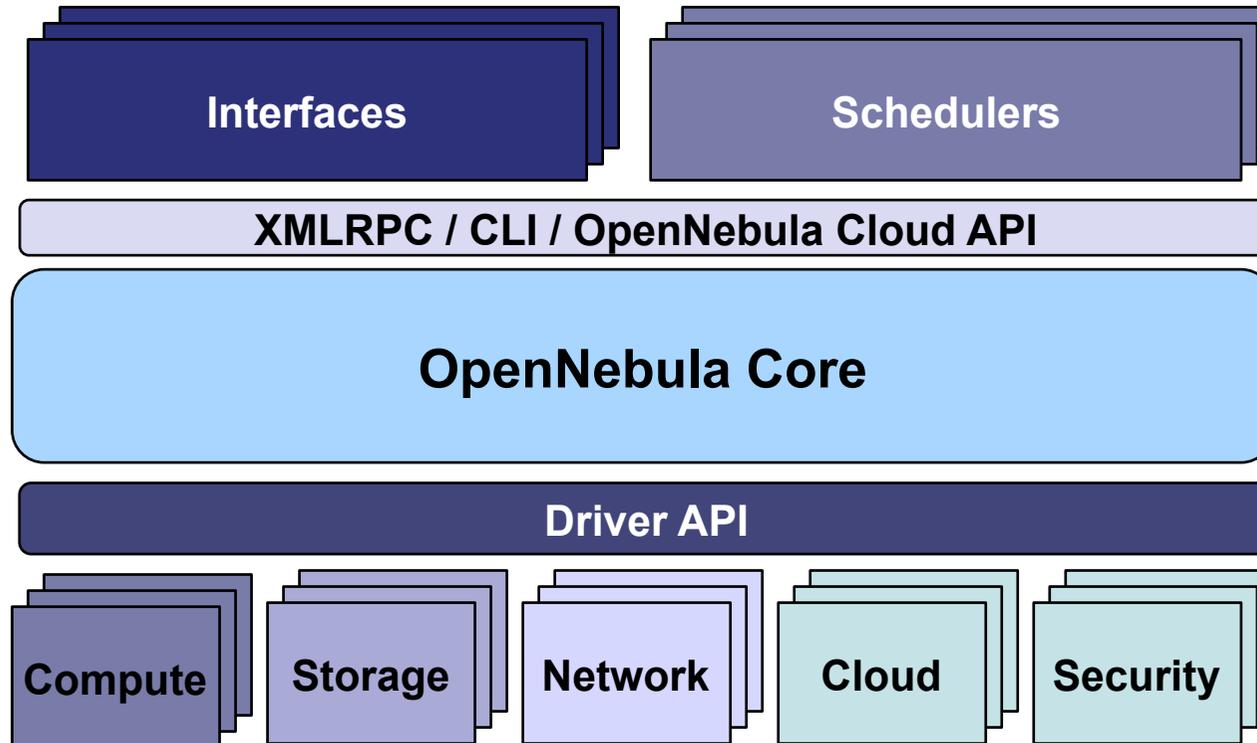
- **Most innovative and advanced solution** to build private, public, federated and hybrid clouds
- **Open and flexible tool** to fit into any datacenter and integrate with any ecosystem component; and **open-source** released under Apache v2.0
- **Efficient and scalable management** of the cloud
- Active and engaged **open community and ecosystem**



Building a Cloud: Open, Flexible, and Extensible

Innovation in Cloud Computing Architectures

A Highly Modular Architecture to Fit into any Existing Datacenter

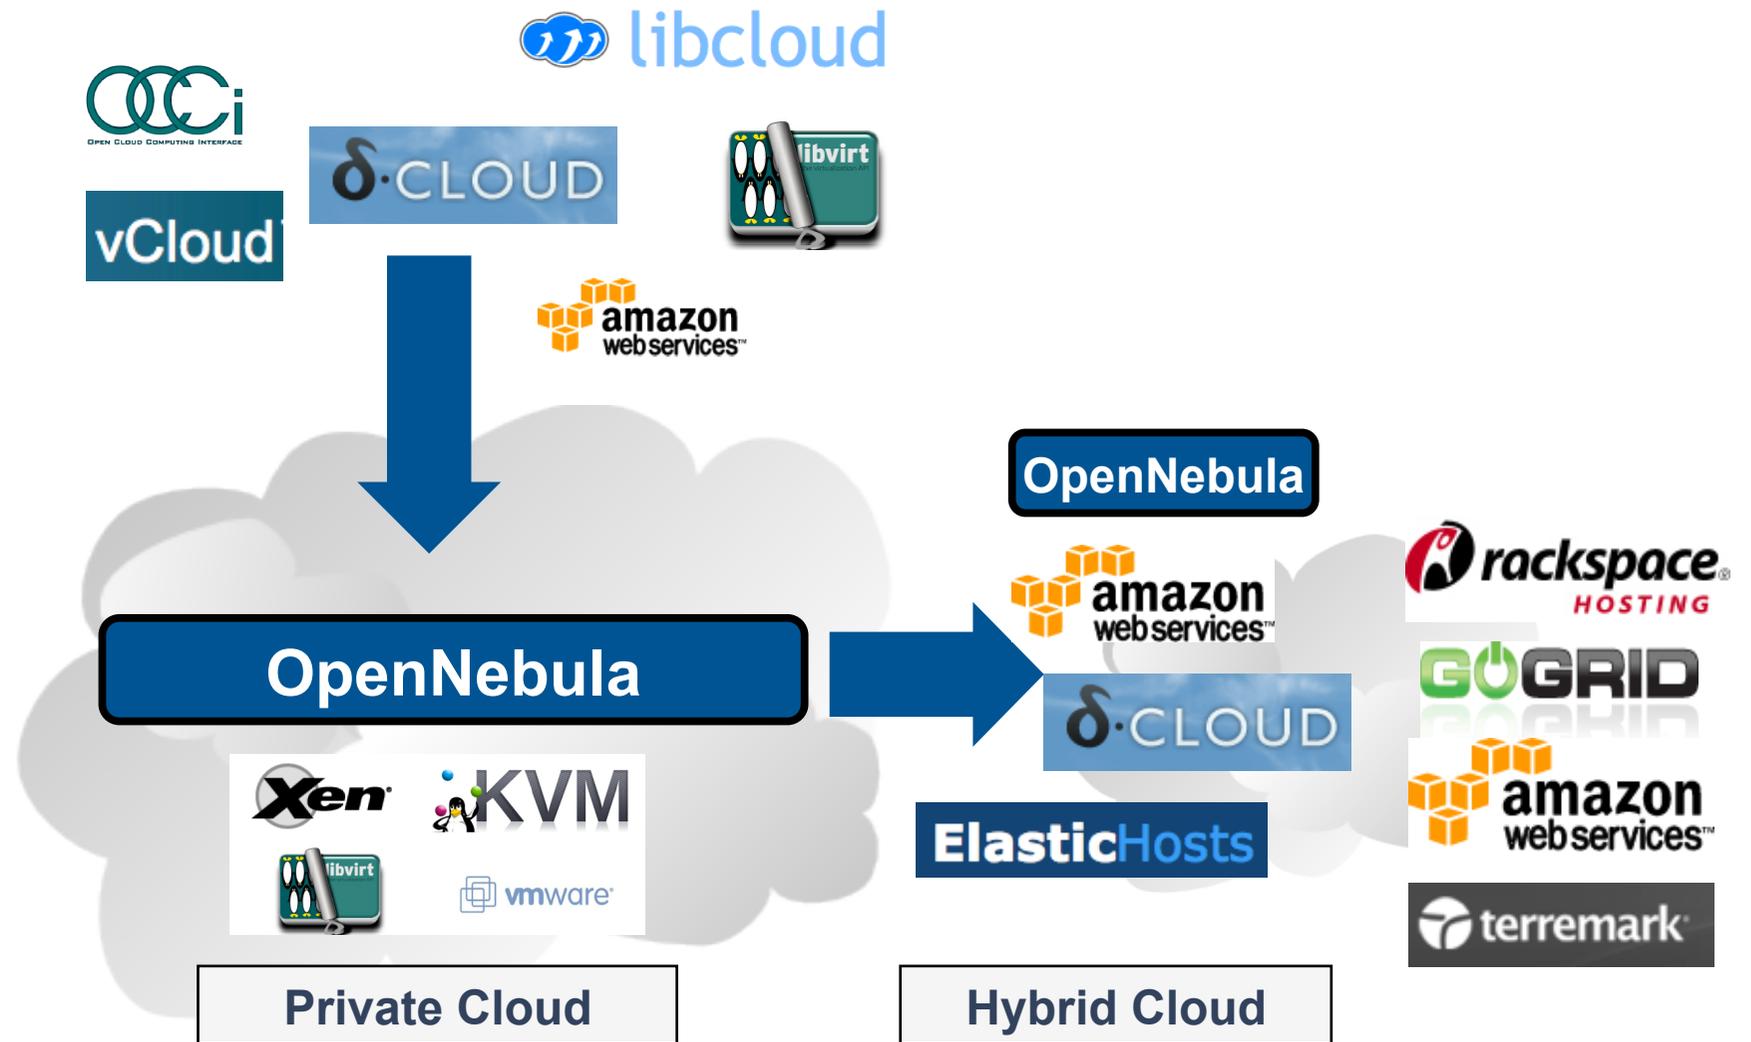




Building a Cloud: Interoperability Map

Innovation in Cloud Computing Architectures

dsa-research.org





Building a Cloud: Experiences in Computing

Innovation in Cloud Computing Architectures

Different Levels of Use: From Experimental to Production

dsa-research.org



Building a Cloud: Experiences in Computing

Innovation in Cloud Computing Architectures

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

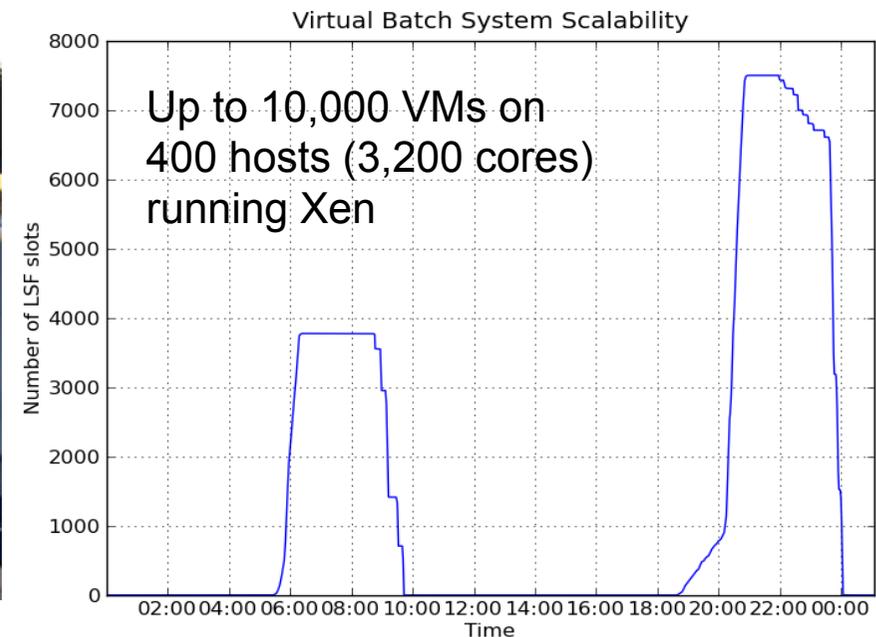
Deployment Cases: Public HPC Cloud



- 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

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



A Tool for Innovation

Innovation in Cloud Computing Architectures

European Projects on Cloud Computing Infrastructures



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



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



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



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

A Tool for Innovation: Enhancing Grid with Cloud

Innovation in Cloud Computing Architectures

 **StratusLab.eu**

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**

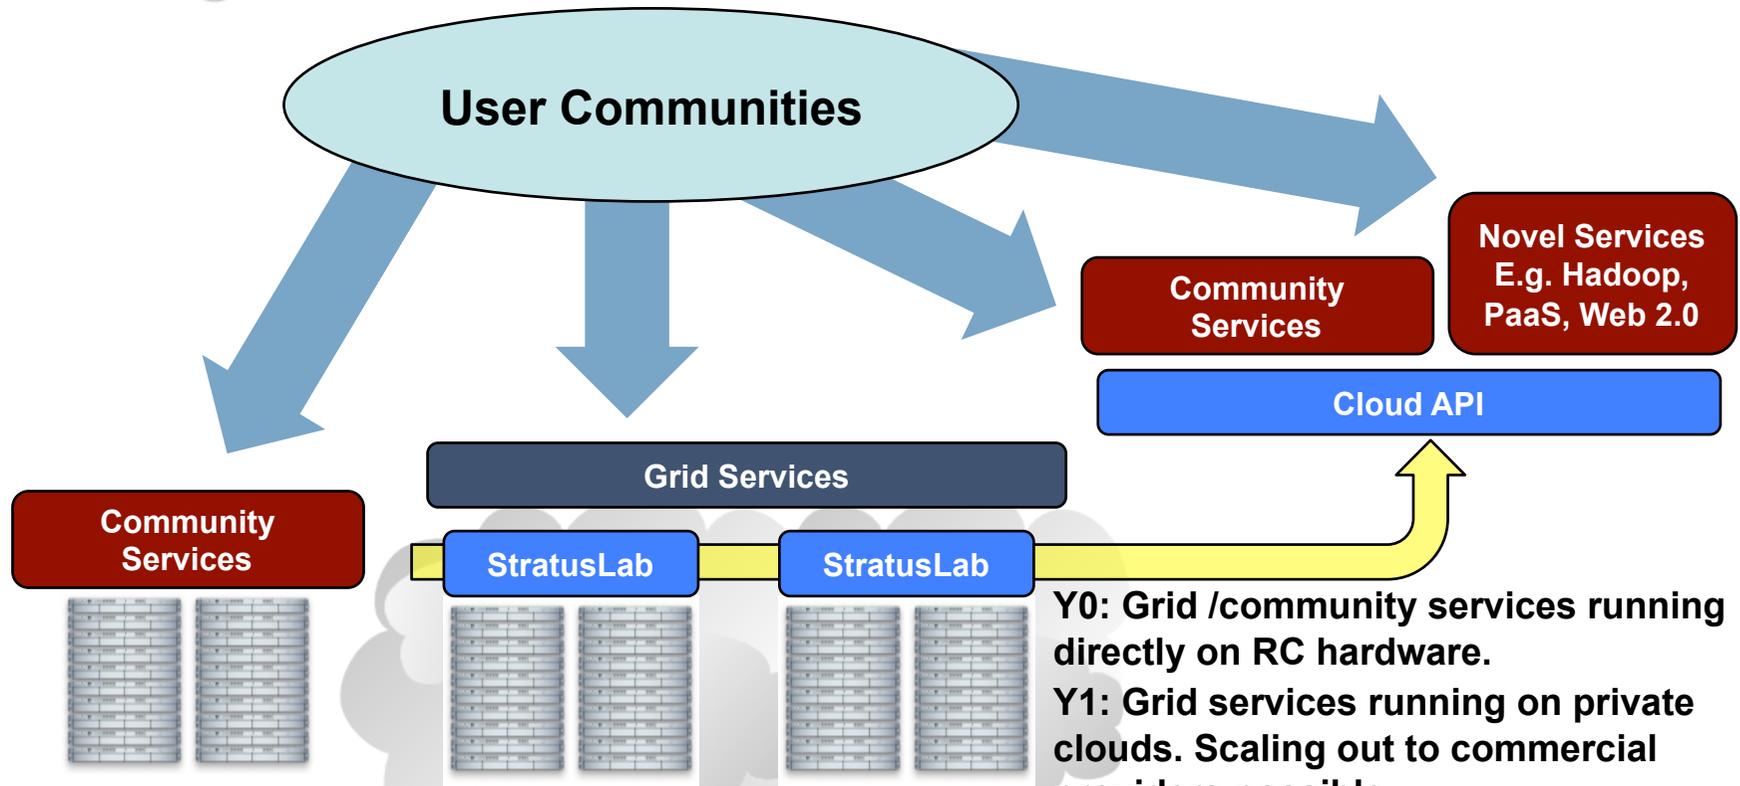


A Tool for Innovation: Enhancing Grid with Cloud

Innovation in Cloud Computing Architectures



dsa-research.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.





Outlook – The Data Center is the Computer

Innovation in Cloud Computing Architectures

IT Resources will be the Next Utility

- **Future enterprise datacenters will look like private Clouds** supporting a flexible and agile execution of virtualized services, and combining local with public Cloud-based infrastructure to enable highly scalable hosting environments
- **Growing number of domain specific and regional Cloud providers implementing a utility computing business model** by offering pay per use resources on-demand
- **Public Clouds will be supported by a network of geographically distributed datacenters** for high availability, end-user service proximity, legal and policy issues...
- **Public Clouds will be interconnected to meet fluctuating demands**
- **Grid sites will offer infrastructure cloud-like interfaces** to address the new resource access demands from the community



Outlook – Research and Technology Challenges

Innovation in Cloud Computing Architectures

Development of Elastic and Scalable Applications

- **Identification of limitations and bottlenecks** of business and scientific applications in Cloud environments
- **Performance and reliability** of business and scientific applications in Clouds
- **Novel applications** of Cloud Computing
- **Grid, HPC and data-intensive computing** in Clouds

Virtual Infrastructure Management and Enablement

- **Novel architectural models** for private, public and hybrid Cloud infrastructures
- **Federation, interoperability and portability** between Cloud providers
- **Metering, monitoring and pricing models** for cloud computing
- **Automatic management of elasticity**
- **Scalable management of physical resources and groups of VMs**
- **QoS and resource allocation**
- Local and cross-site placement optimization algorithms for **energy efficiency**, load balancing, and **high availability and SLA commitment**.
- **Advance reservation of capacity**



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
- **New EU Projects** (StratusLab, BonFIRE, 4CaaSt) provide **funding until 2013**

Other Sponsors

- **C12G LABS** • **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!



Get Involved in the OpenNebula Community!

Build Your Cloud!

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

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

More Information

More info, downloads, mailing lists at



The screenshot shows the OpenNebula.org website. The header includes the logo "OpenNebula.org" and the tagline "The Open Source Toolkit for Cloud Computing". Below the header is a navigation menu with links for Home, About, Documentation, Software, Support, Community, Cloud, Dev, and Blog. The main content area is divided into several sections: "The Leading and Most Advanced Solution for Cloud Computing" with a list of supported environments (Xen, KVM, VMware, Amazon EC2, ElasticHosts, EC2 Query, OGF OCCl, vCloud); "Getting Started" with a three-step process (Download OpenNebula, Read the Documentation, Engage the Community); "The Fastest Track to Cloud Computing with C12G OpenNebula Express" with a button "Your Cloud on Ubuntu and RHEL/CentOS in Few Minutes"; "Featured Quotes" with a quote from SurfNet; and "Announcements" with a list of recent news items.

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)