Workshop GARR - Calcolo e Storage Distribuito
Rome, Italy, November 30th, 2012

Federated Cloud Computing
The OpenNebula Experience

Ignacio M. Llorente

Project Director

OpenNebula.org

Acknowledgments

The research leading to these results has received funding from the *Ministerio de Ciencia e Innovación* of Spain through research grant TIN2009-07146.

This presentation is about:

- Infrastructure as a Service Cloud Computing
- Private Clouds in Science and HPC
- Federated Cloud Computing
- Cloud federation in Grid infrastructures

What is Cloud Computing?

Provision of IT Capabilities as a Service

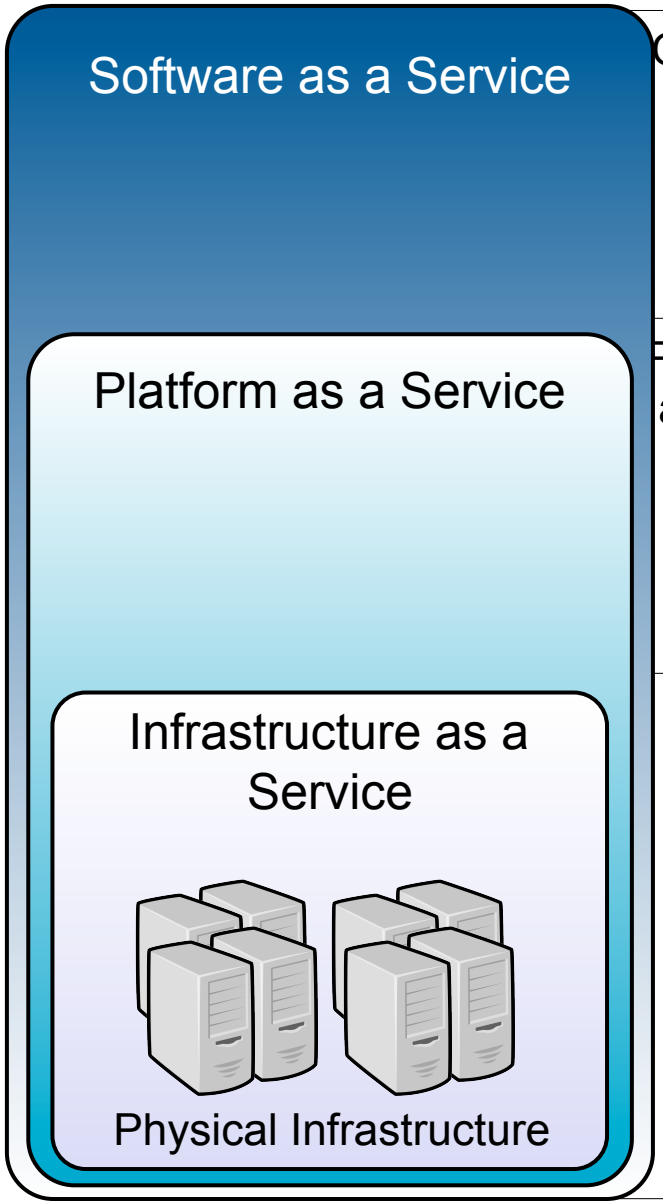

What	Who
On-demand access to any application	End-user (does not care about hw or sw)
Platform for building and delivering web applications	Developer (no managing of the underlying hw & swlayers)
Raw computer infrastructure	System Administrator (complete management of the computer infrastructure)

What is Cloud Computing?

IaaS Clouds for Provision of Virtualized Resources as a Service

- Management Console
- Simple REST API's

- Raw infrastructure resources

Manage Instances

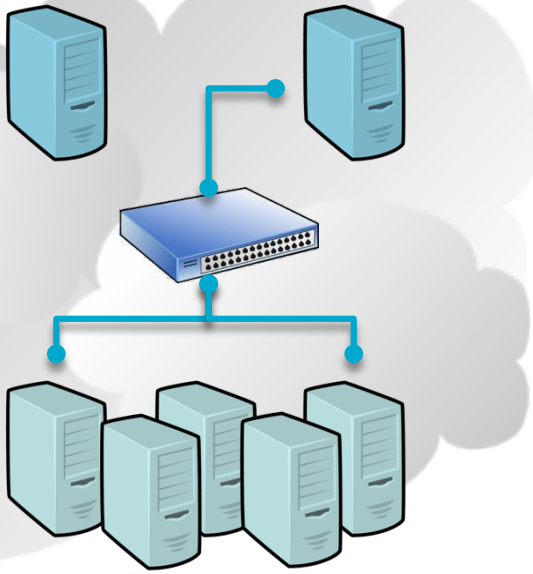

- Pay-as-you-go & elastic capacity

Different Models of Deployment

Model	Infrastructure	Cloud Cases
Private	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	Owned by a single organization and made available to other organizations over the Internet	• Commercial cloud providers, mostly hosting providers, with limited control/security • Science public clouds by ICT service centers to enable scientific projects or experiment with cloud computing
Virtual Private	Owned by a single organization and made available to other organization over a dedicated private network	• Telecom cloud providers with premium solutions with additional control/security
Hybrid	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

Open Cloud Management Solution for Building and Managing Virtualized Data Centers

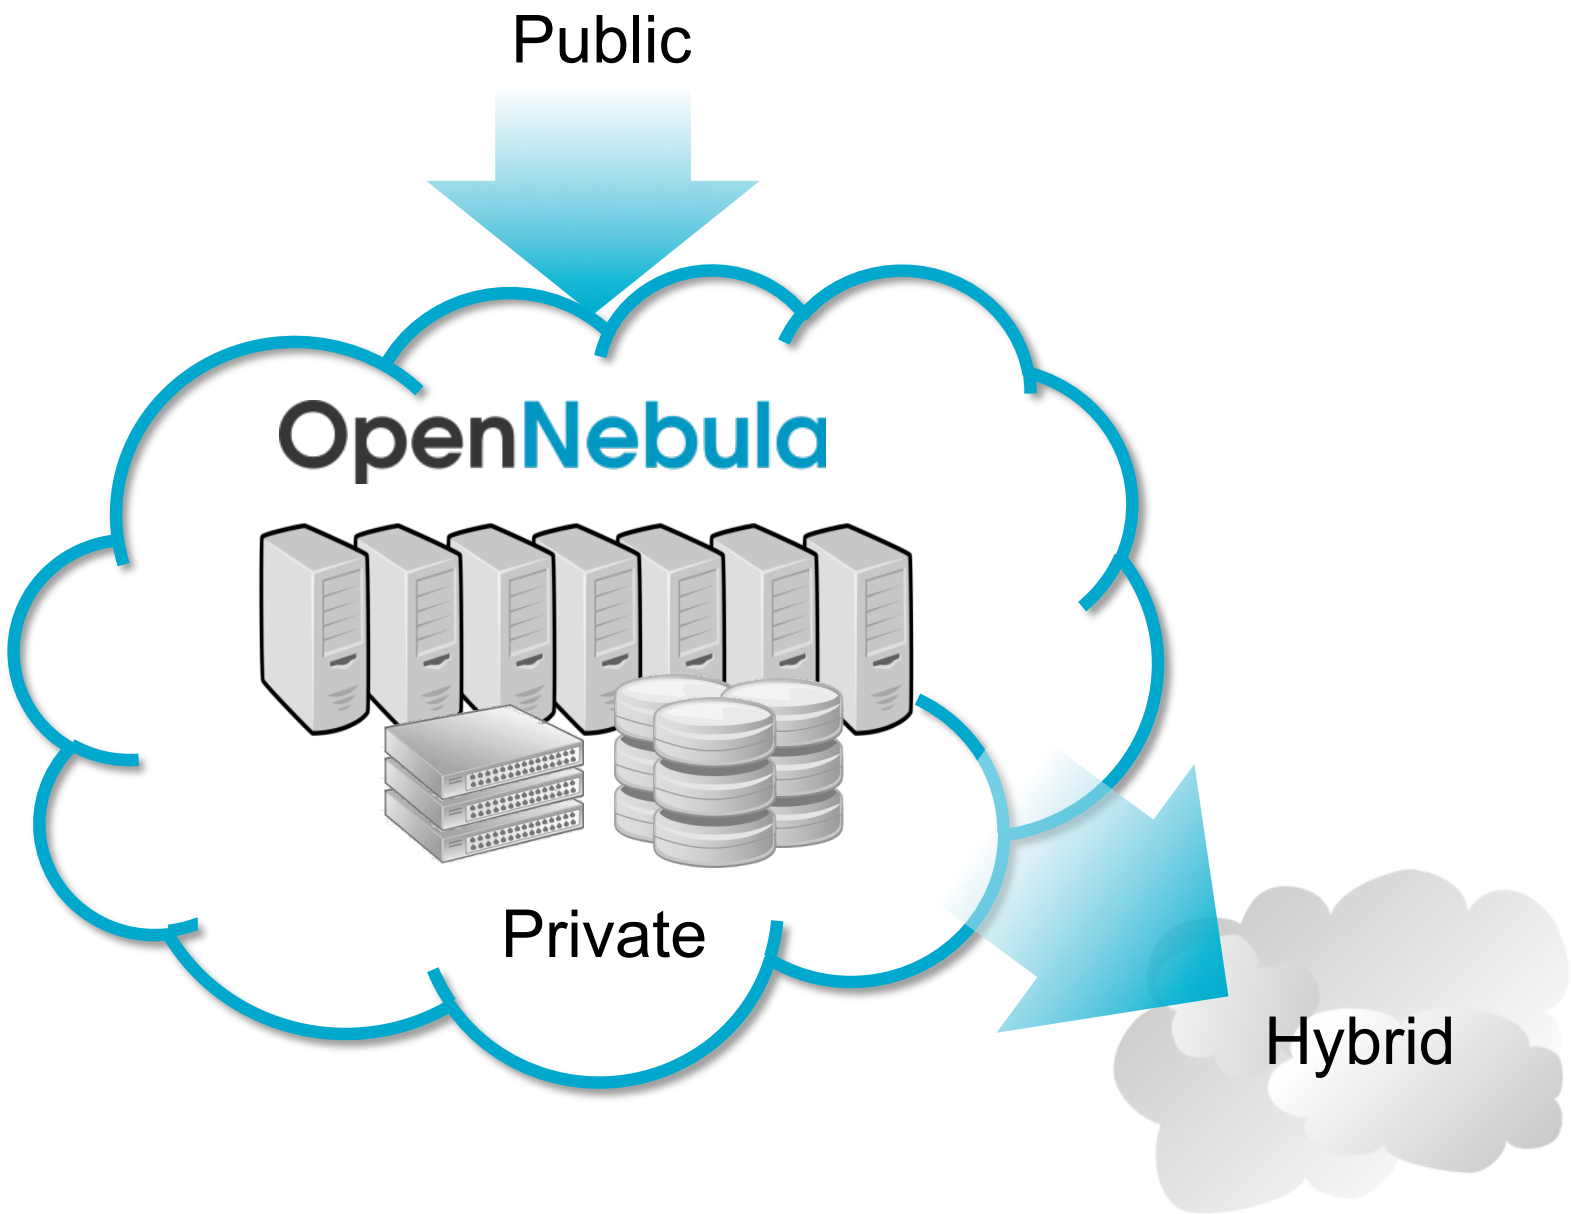

What is OpenNebula?

Most Advanced and Flexible, Enterprise-grade IaaS Cloud Manager

Adaptable
- Customizable and Extensible

Proven
- Many Massive Scale Production Deployments

Powerful
- Most Advanced Enterprise-class Functionality

No Lock-in
- Platform Independent and Interoperable

Interoperable
- Most popular cloud APIs and standard based

Openness
- Fully open-source, Apache license

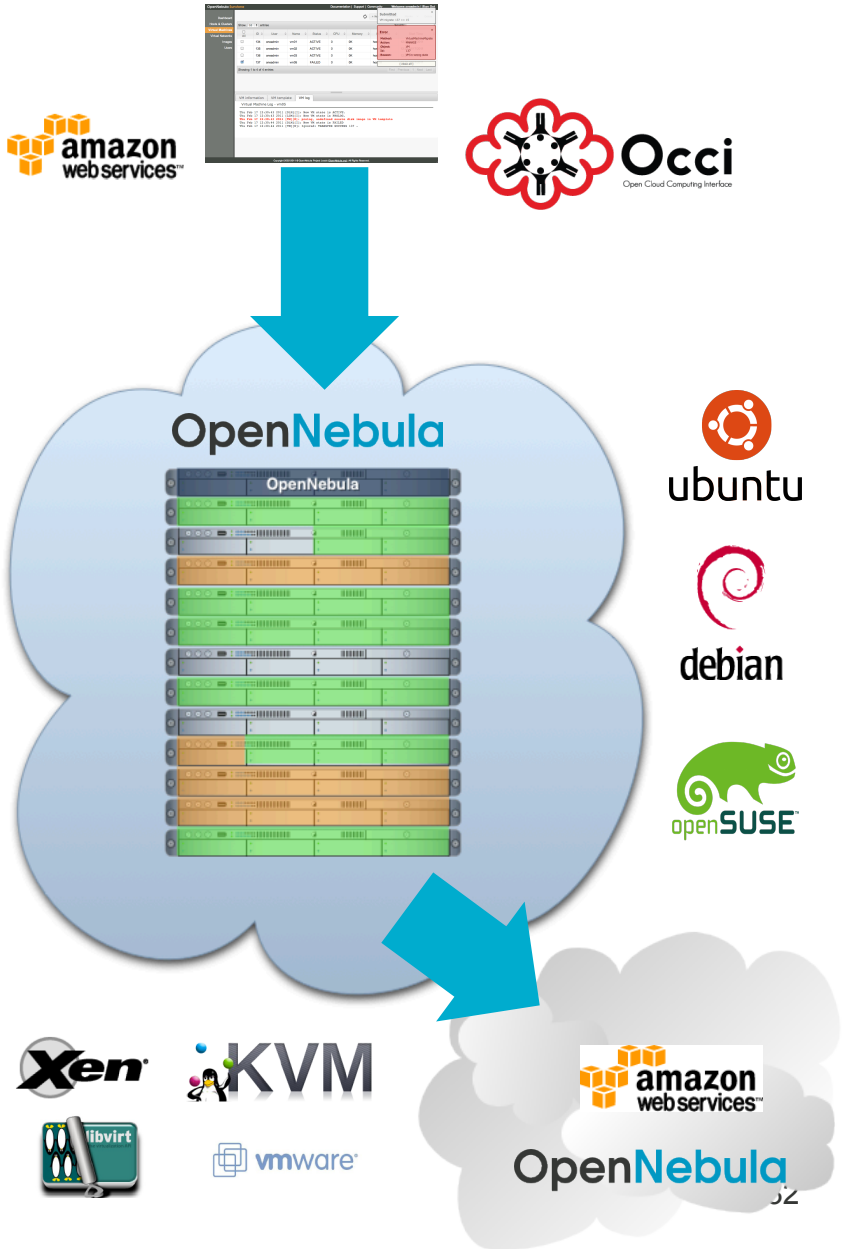

What is OpenNebula?

User-driven Open Platform for Innovation

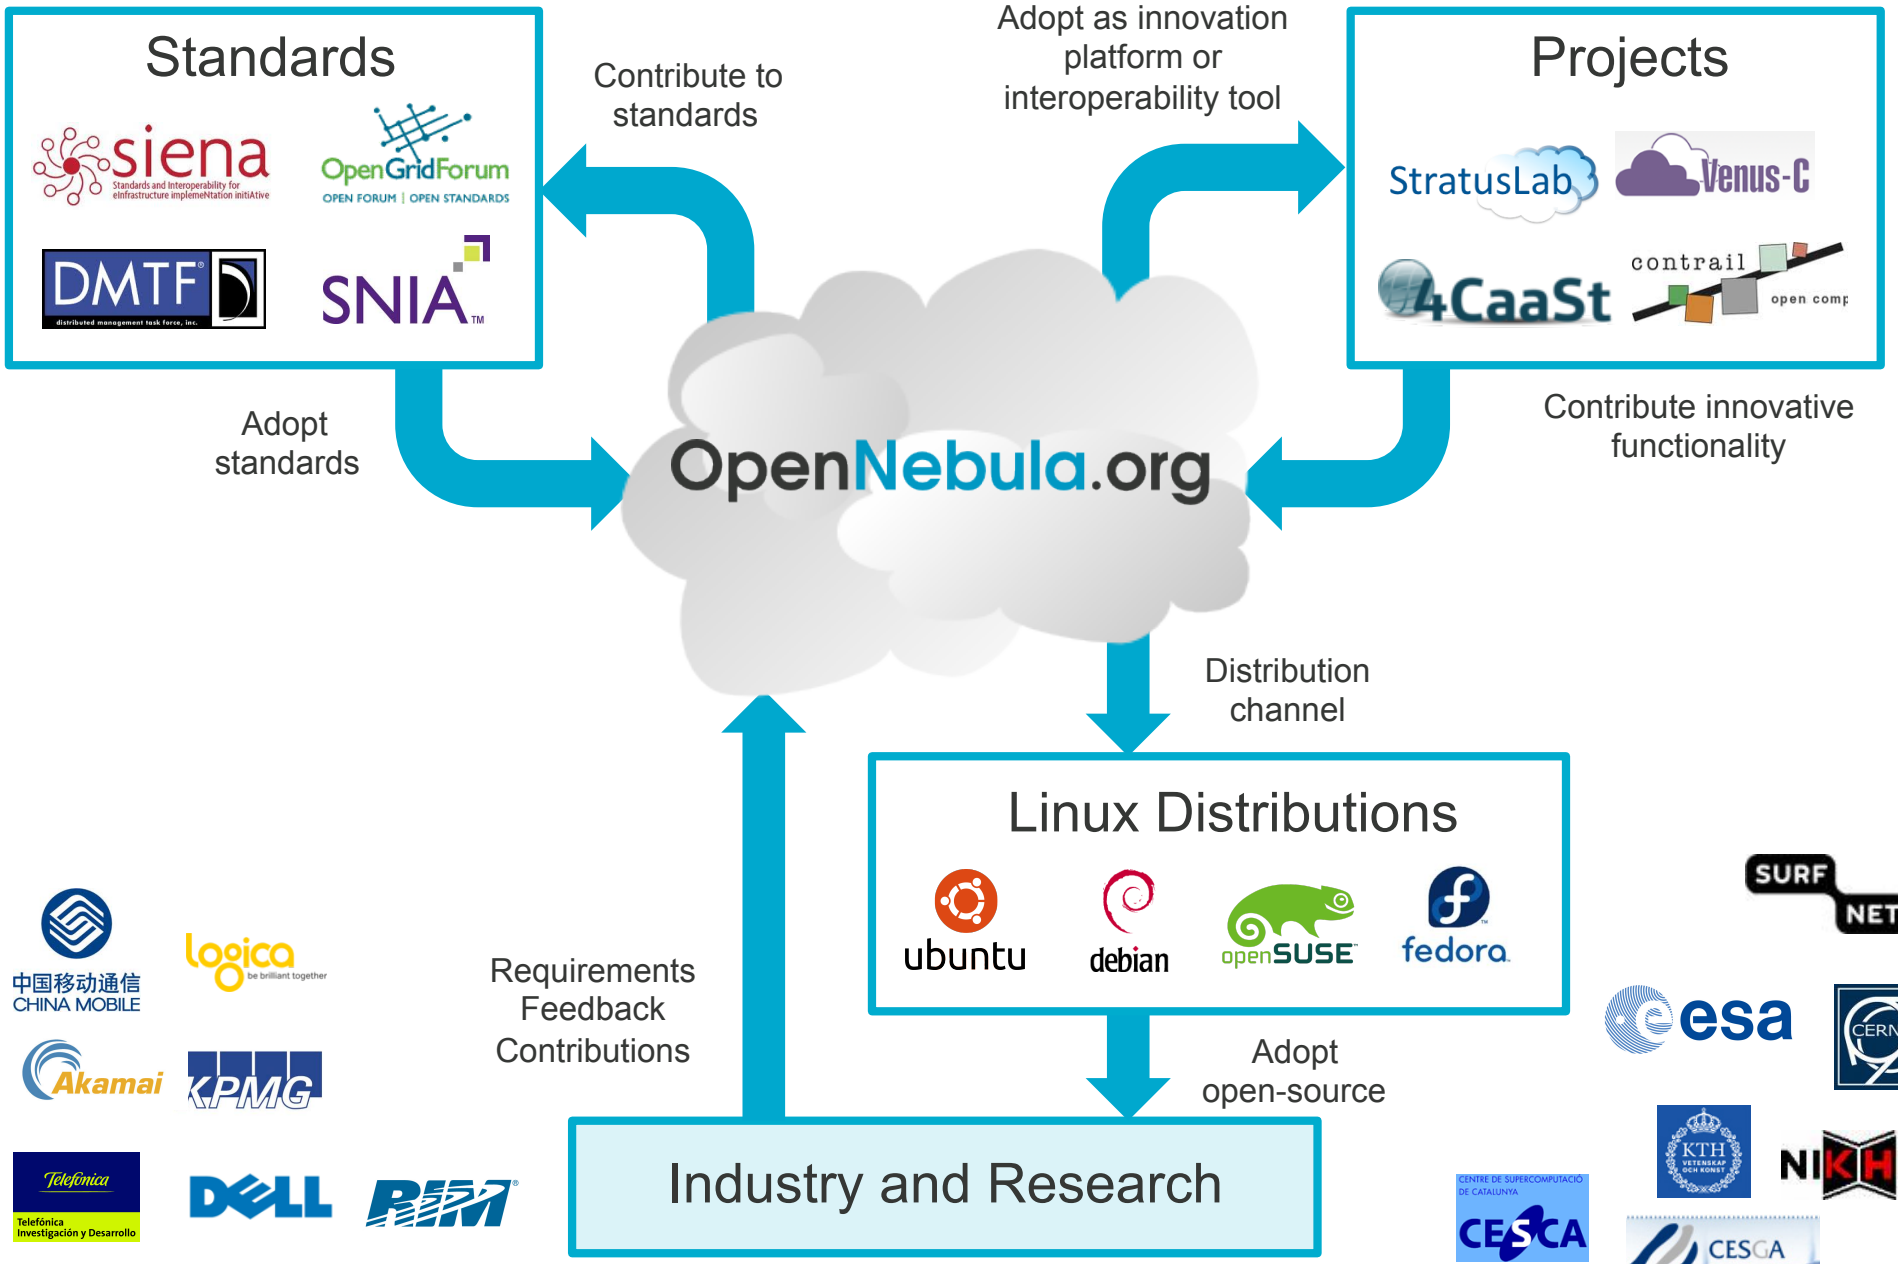

The Pre-cloud Era

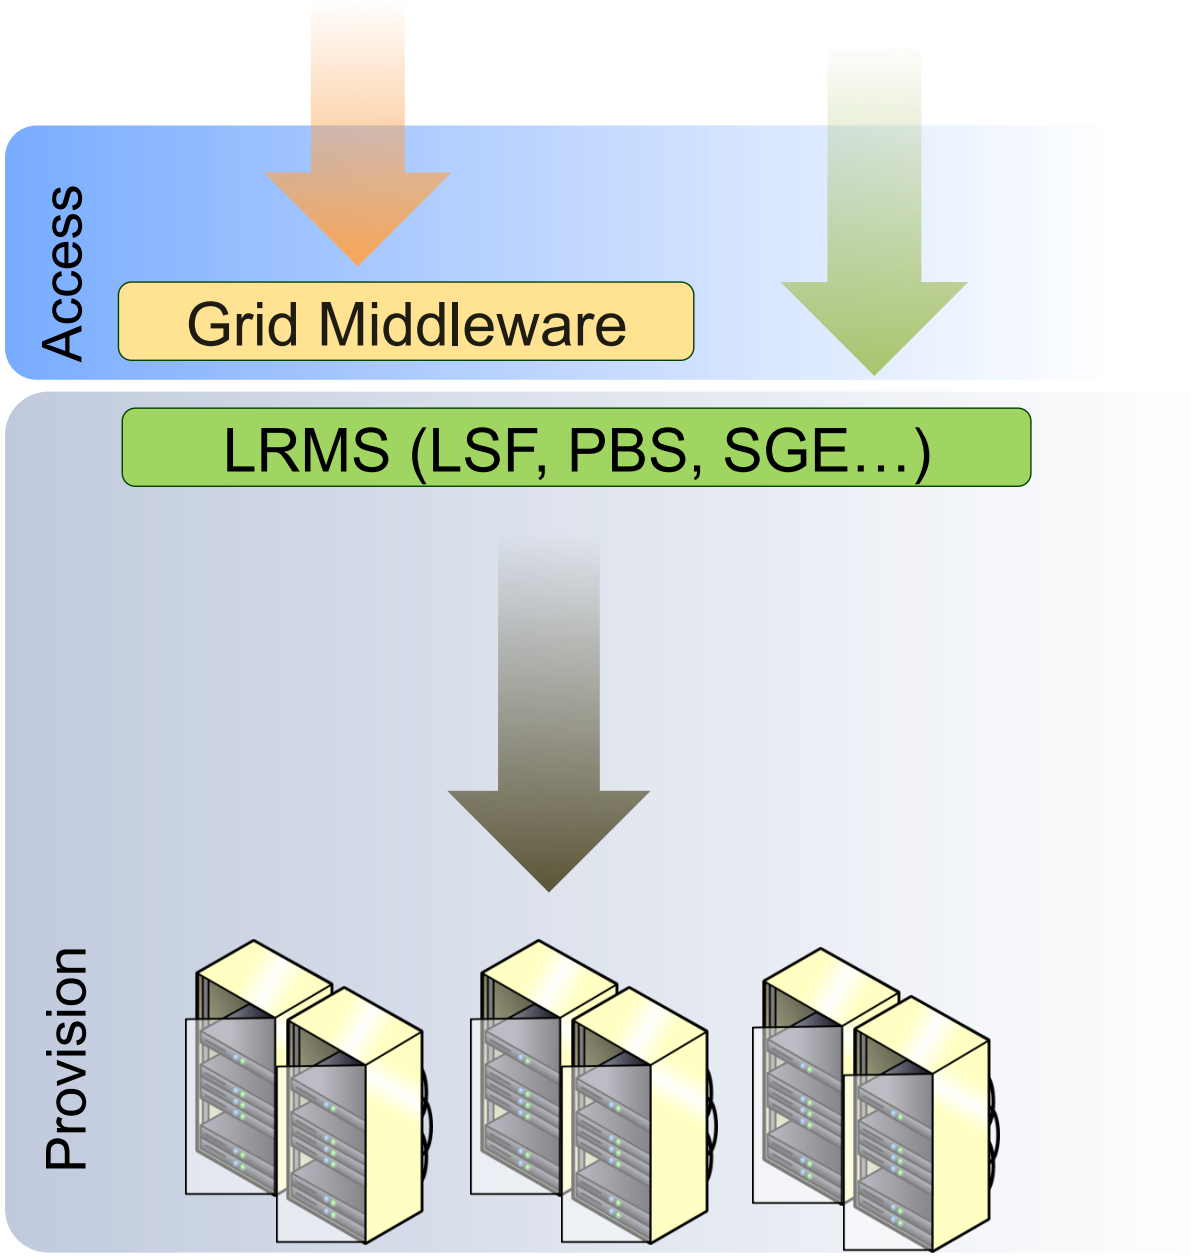

OpenNebula as an Infrastructure Tool

Service/Provisioning Decoupling

Access

Grid Middleware

- Common interfaces
- Grid integration

Service

LRMS (LSF, PBS, SGE…)

Virtual Worker Nodes

- Custom environments
- Dynamic elasticity

Provision

OpenNebula.org

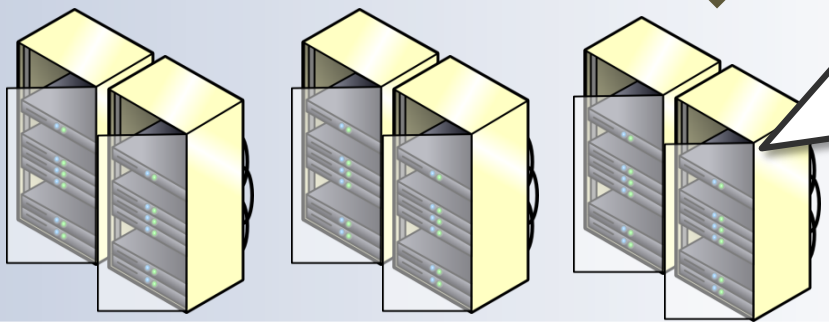

- Consolidation of WNs
- Simplified management
- Physical – Virtual WNs
- Dynamic capacity partitioning
- Faster upgrades

OpenNebula as an Provisioning Tool

Access

IaaS Interface

- Simple Provisioning Interface
- Raw/Appliance VMs

Service

Pilot Jobs, SSH…

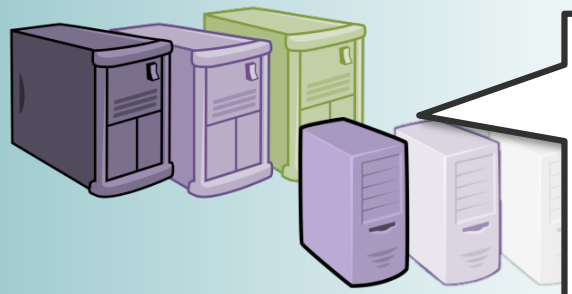

- Dynamic scalable computing
- Custom access to capacity
- Not only batch workloads
- Not only scientific workloads

Provision

OpenNebula.org

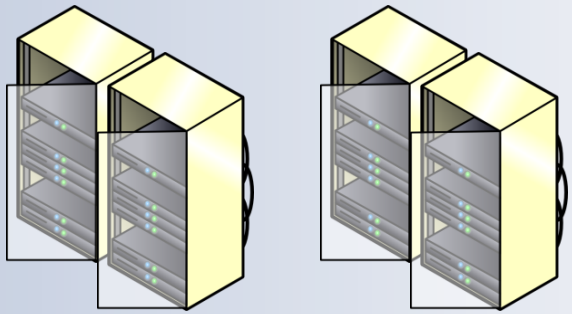

- Improve utilization
- Reduced service management
- Cost efficiency

Examples: CERN's lxcloud

http://blog.opennebula.org/?p=620

Goal

- Virtualized HTC Batch Nodes: limited life-time, dynamic capacity
- EC2 Query: CernVM

Deployment Notes

- Custom network integration MAC/IP pinning
- Fast image distribution & boot: BitTorrent + LVM snapshots
- OpenNebula tests up to 20,000 VMs (great feedback!)

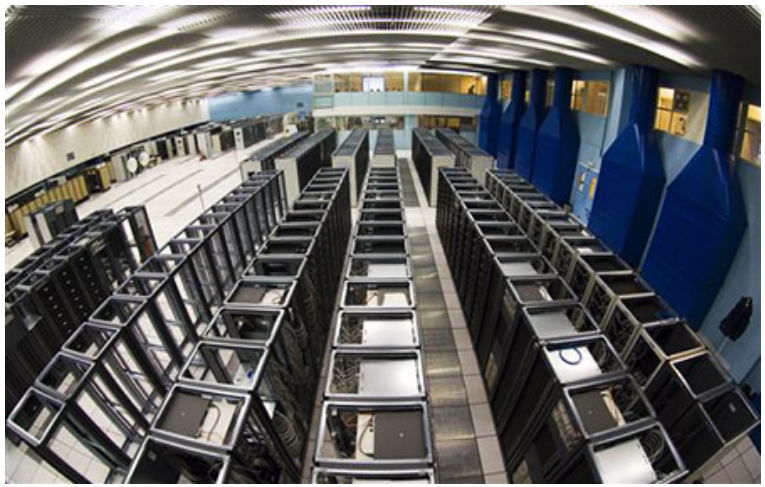

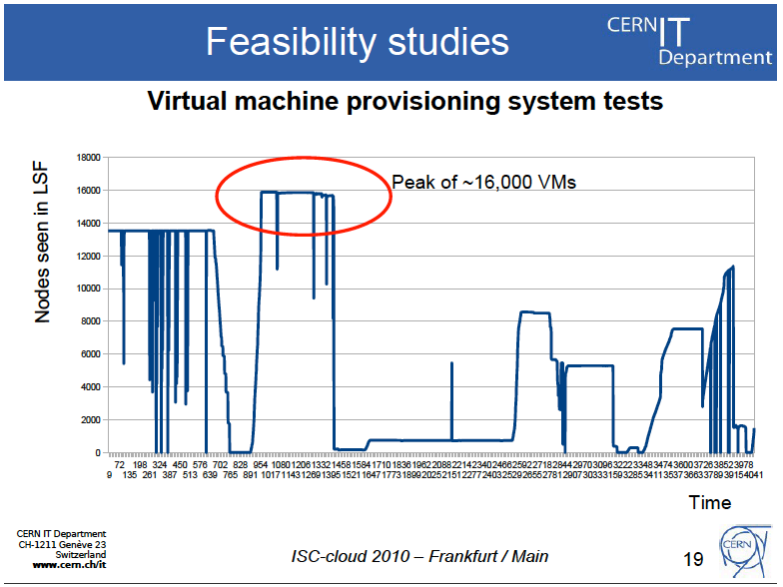

Examples: FermiCloud

http://www-fermicloud.fnal.gov/

Goal

- Scientific stakeholders get access to on-demand VMs
- Developers & integrators of new Grid applications
- MPI and legacy applications

Deployment Notes

- VMs access Fermilab Networking and Storage Services
- OpenNebula + X509 support (contributed back!)
- Other areas: HA, Batch queues look-ahead, cluster on-demand…

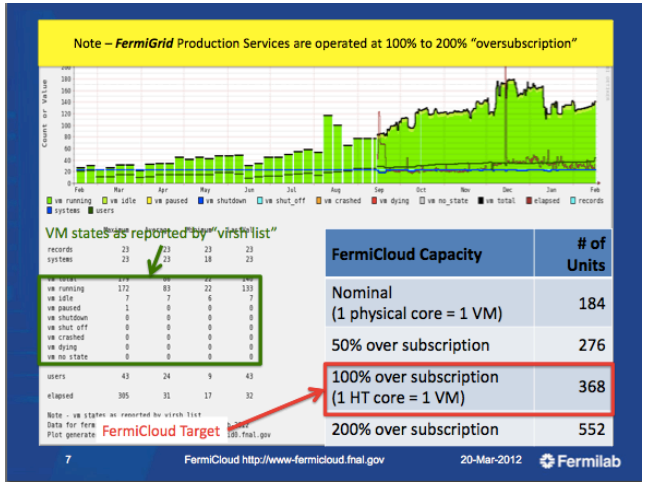

Examples: SARA

https://www.cloud.sara.nl/

Goal

- HPC clusters on-demand: Self service, elastic and fully configurable HPC systems
- Offer production infrastructure for several engineering and scientific communities: Bioinformatics, ecology…

Deployment Notes

- Low latency network for HPC
- Extensions in accounting, GUI, firewalling… (contributed back!)

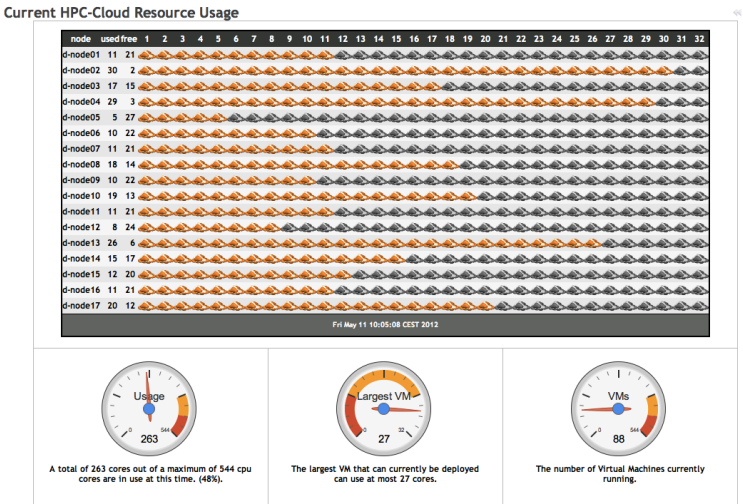

One of Our Main User Communities

Supercomputing Centers

Research Centers

Distributed Computing Infrastructures

Next Step in the Evolution of an Utility

Utility Generation

Utility Distribution

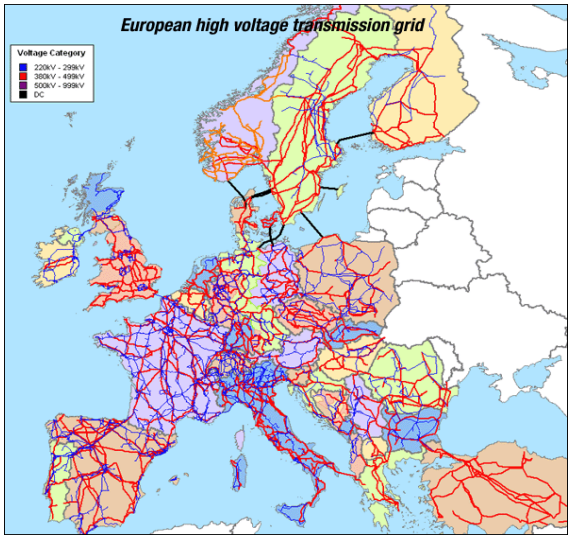

Utility Grid

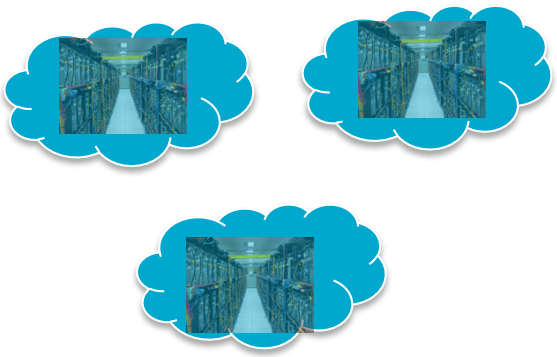

Benefits of Federation

Scalability
- Cloudbursting to address peak demands

Collaboration
- Sharing of infrastructure between partners

Multi-site Deployments
- Infrastructure aggregation across distributed data centers

Reliability
- Fault tolerance architectures across sites

Performance
- Deployment of services closer to end users

Cost
- Dynamic placement to reduce the overall infrastructure cost

Energy Consumption
- Minimize energy consumption

The Cluster Computing Use Case

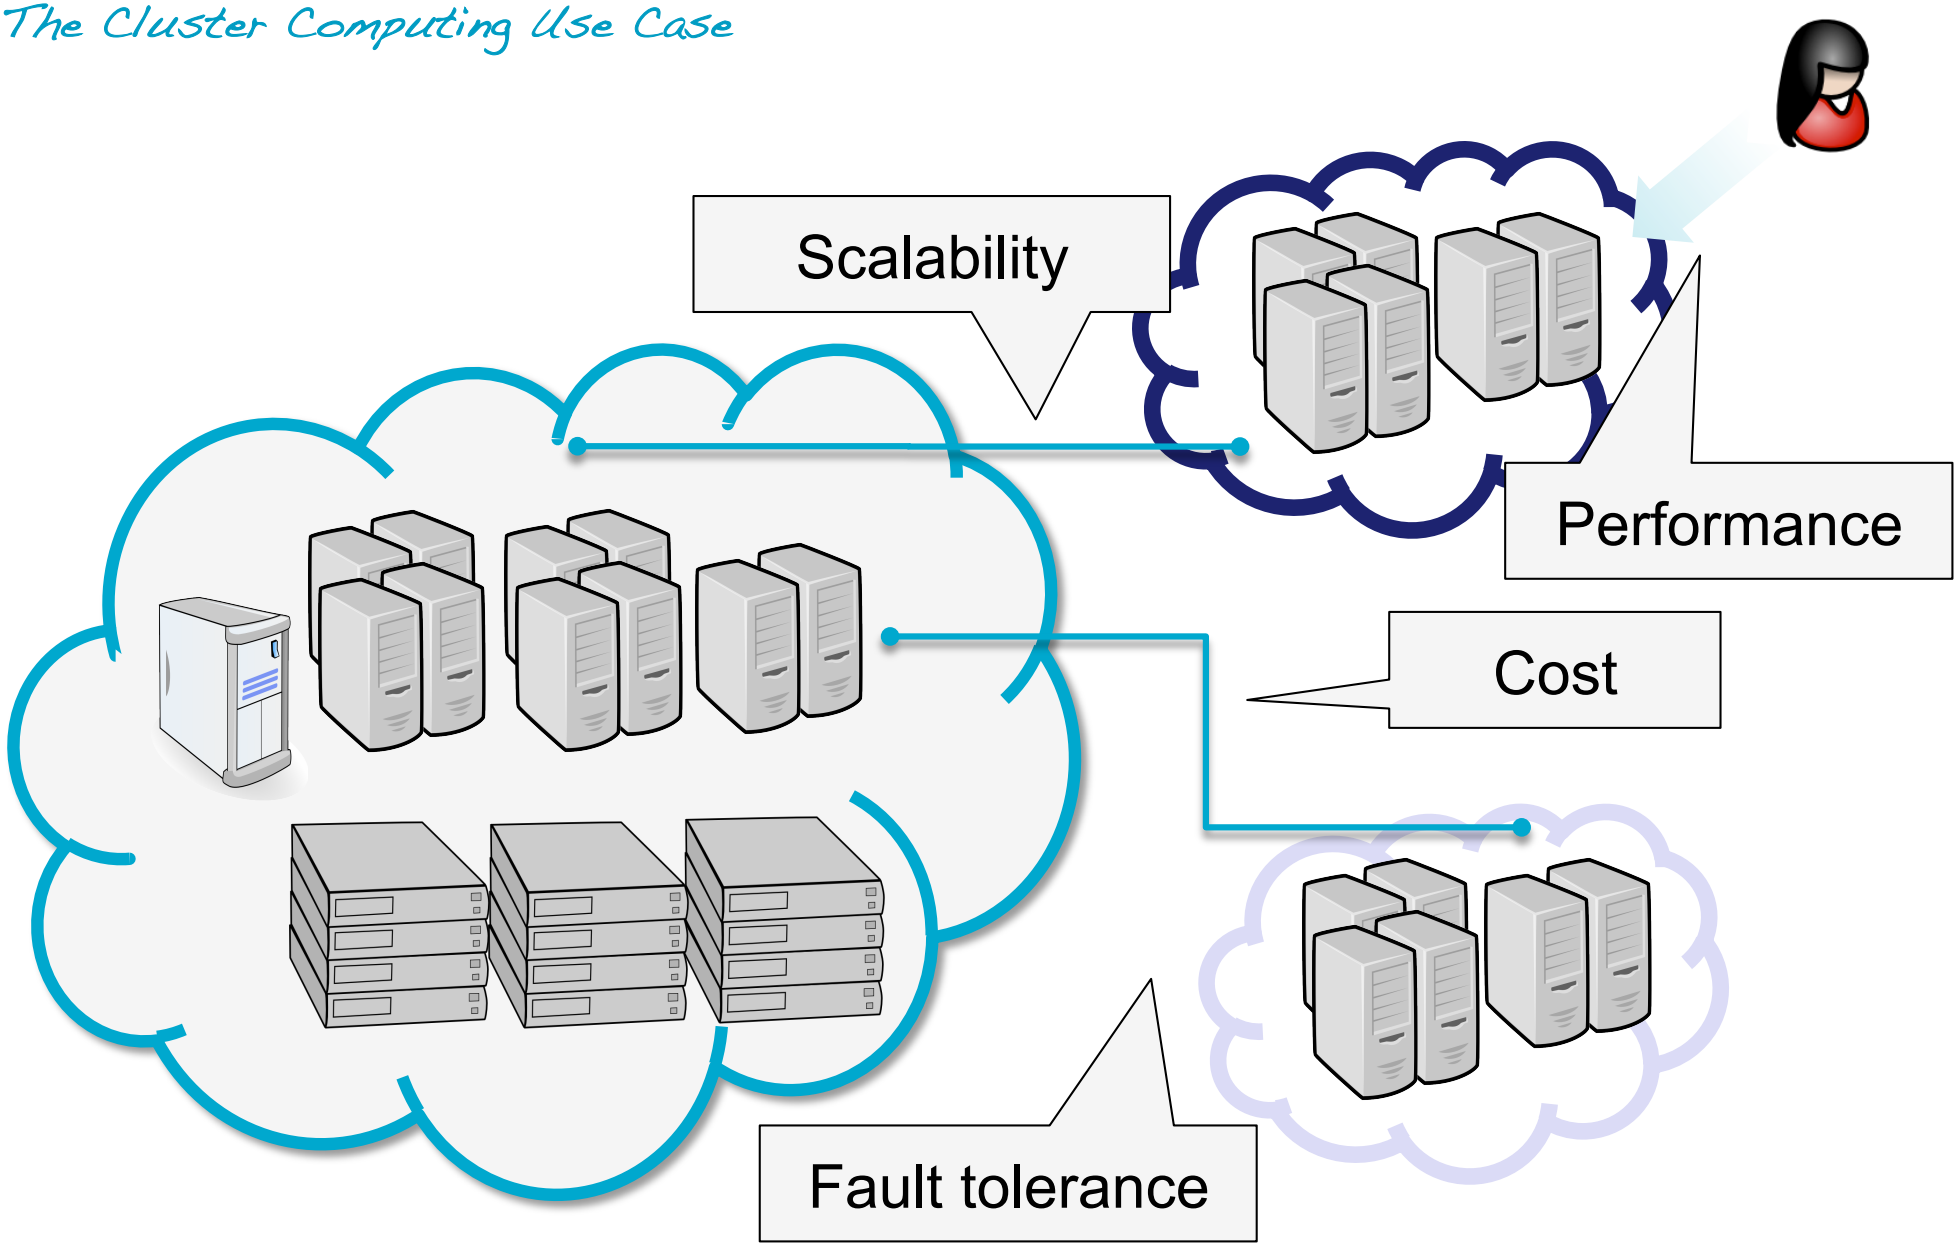

Different Levels of Control, Monitoring, Cross-site Functionality and Security

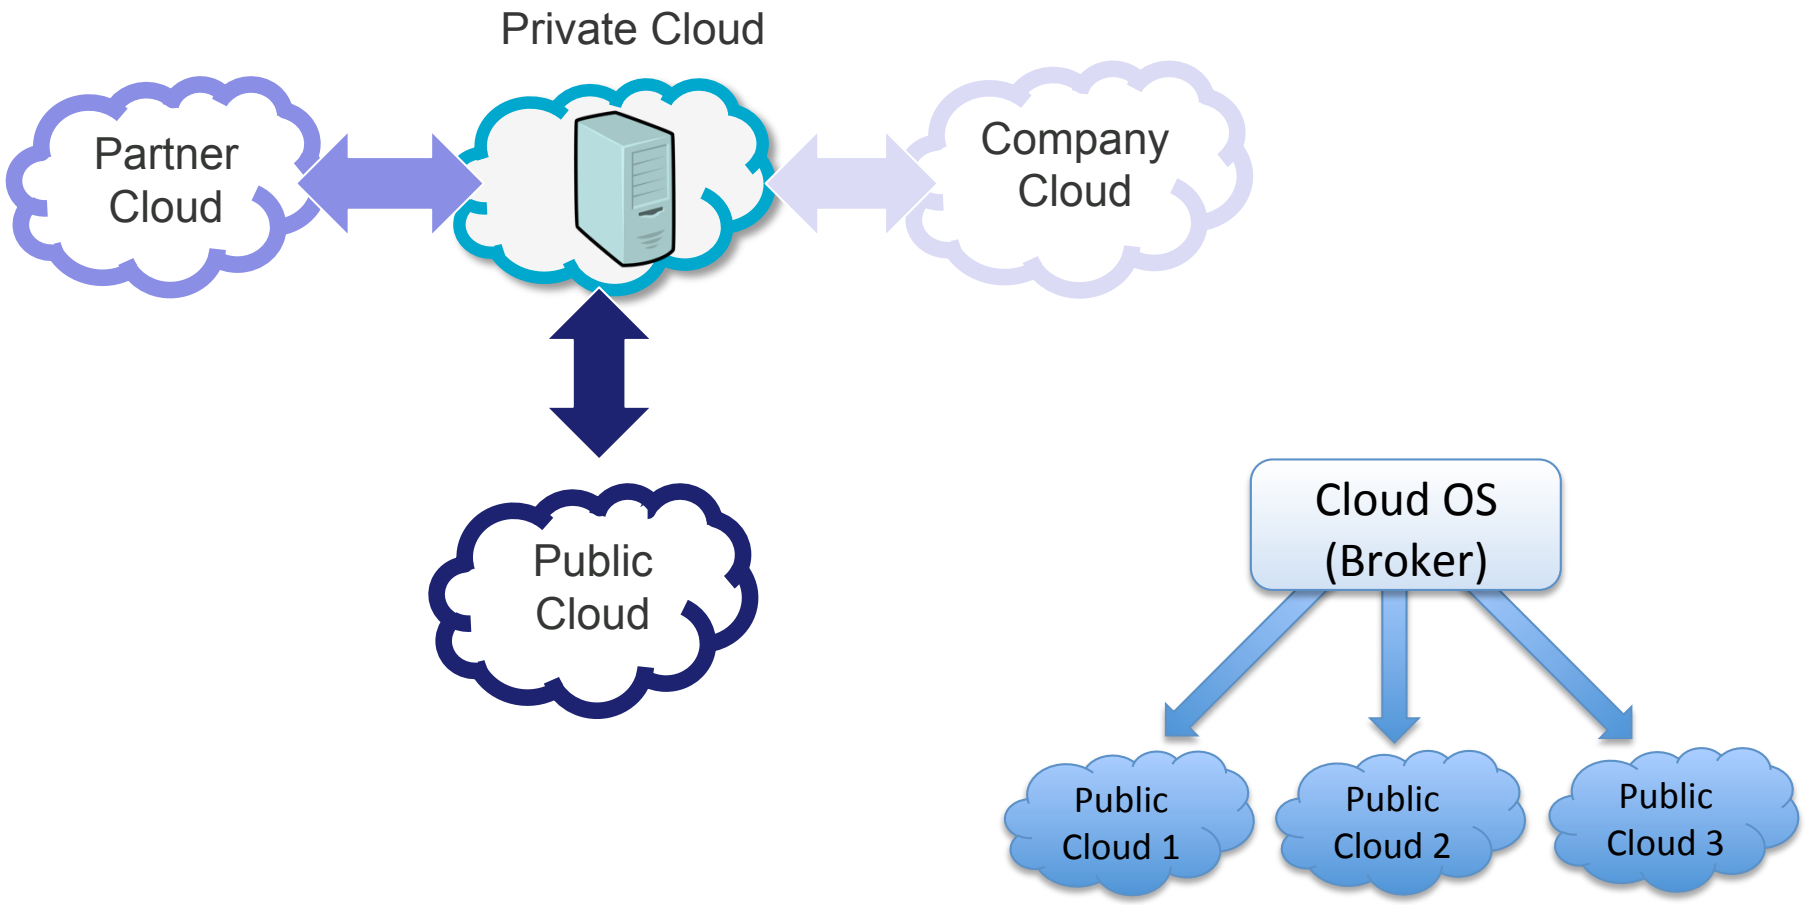

Federation with a Cloud without Interoperation Support

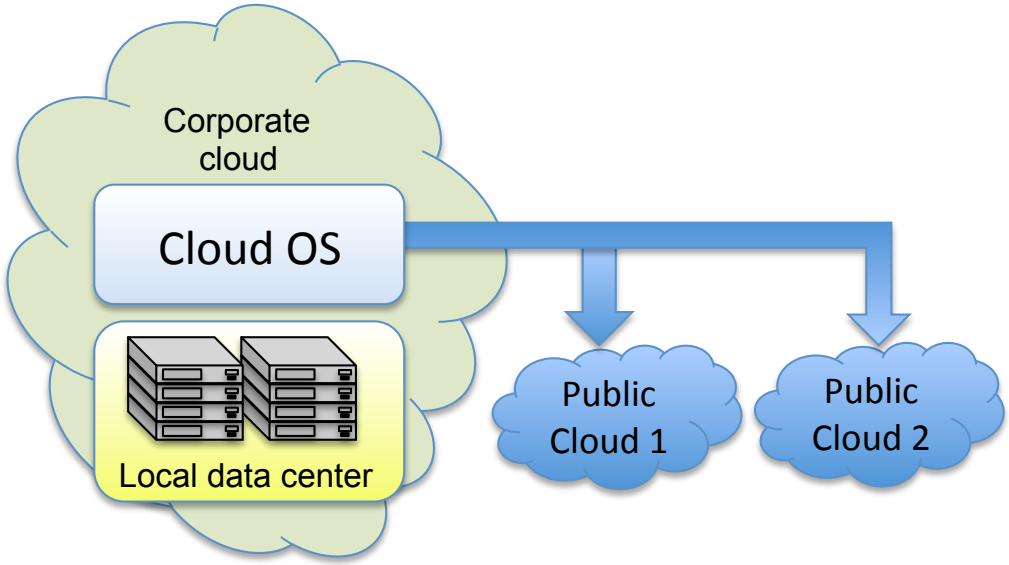

Aim	• Meet peak demands
Control	• Basic operations over VMs (start, shutdown, restart…) • Different instance types
Monitoring & Accounting	• Basic virtual resource monitoring (resource consumption…)
Cross-site	• None
Security	• Single account representing the organization

Dynamic Combination of Local with Remote Cloud Resources

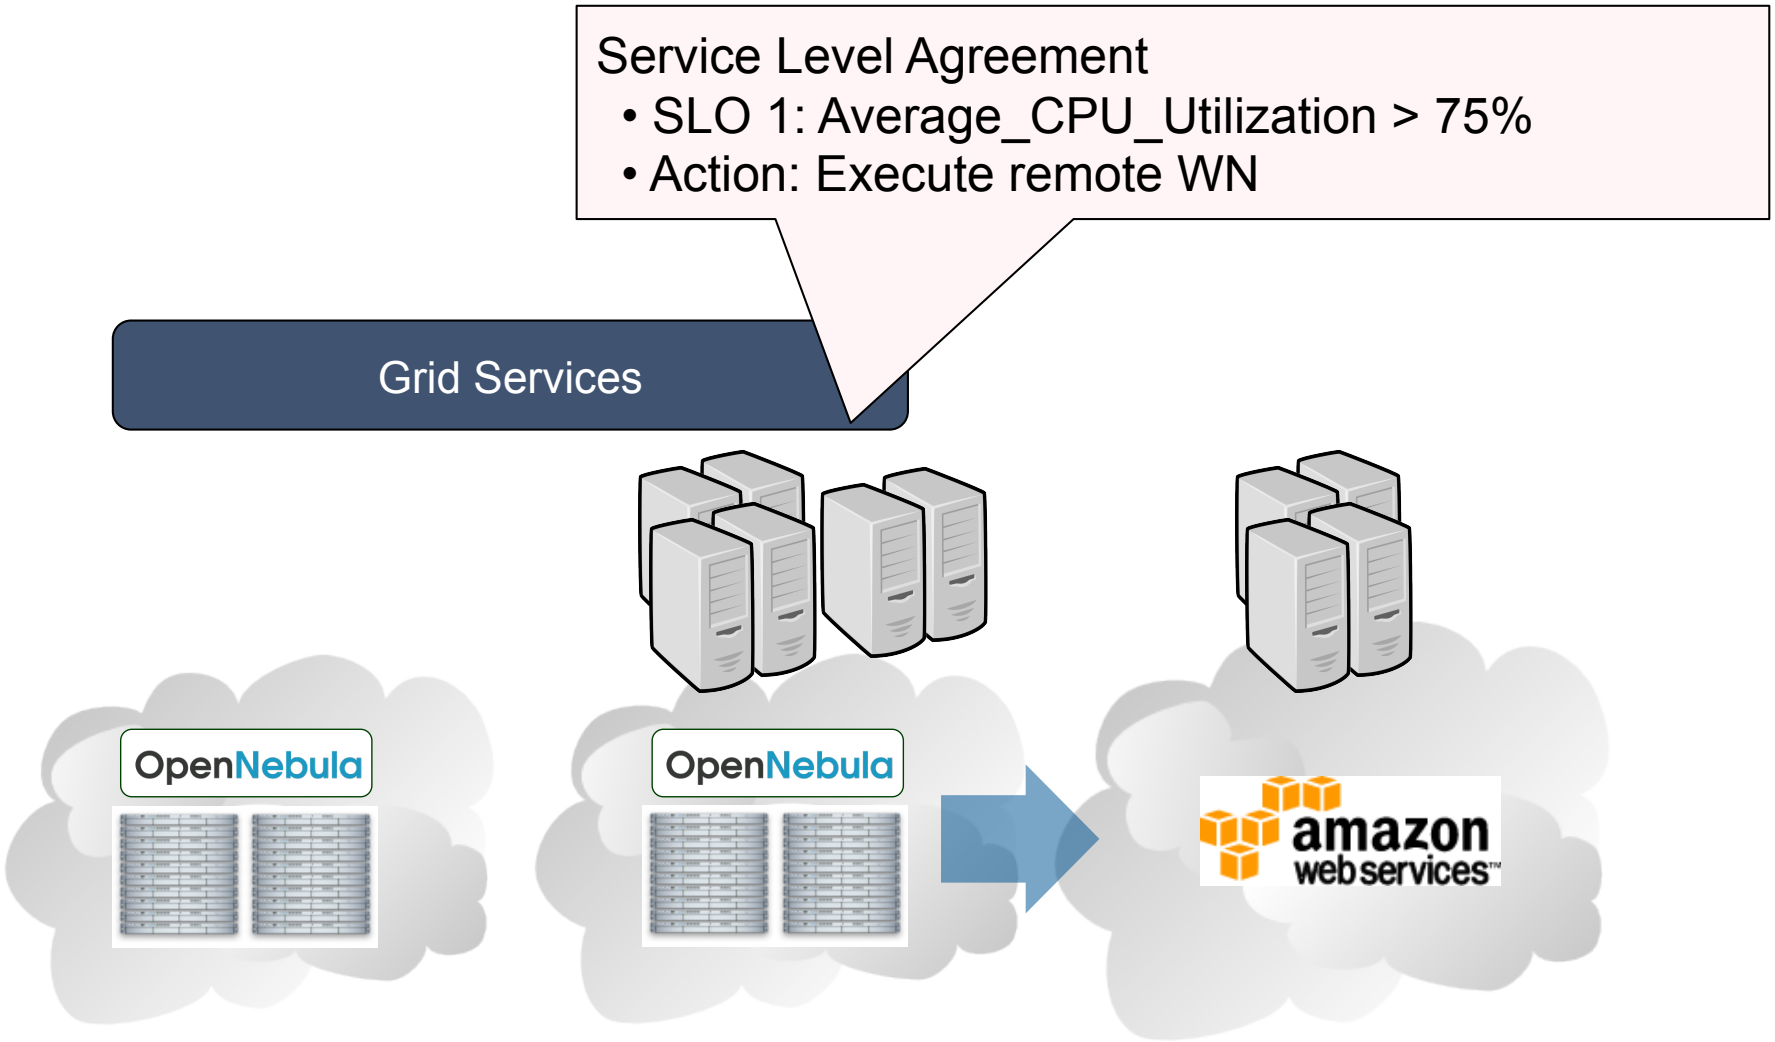

Partially Coupled Federation - Aggregated Cloud Architecture

Federation with a Cloud with Partial Interoperation Support

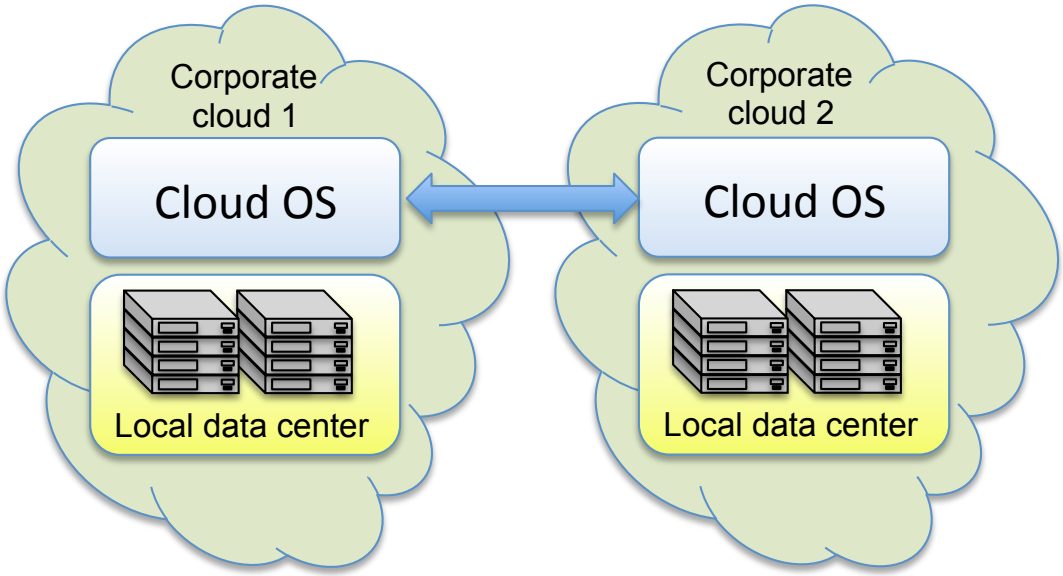

Aim	• Sharing of resources to meet peak demands
Control	• Advanced operations over VMs (live migration…) • VM location and affinity constraints
Monitoring & Accounting	• Advanced virtual resource monitoring (energy consumption, VM placement…)
Cross-site	• Virtual networks • Virtual storage
Security	• Framework agreement

Cloud Broker Architecture

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

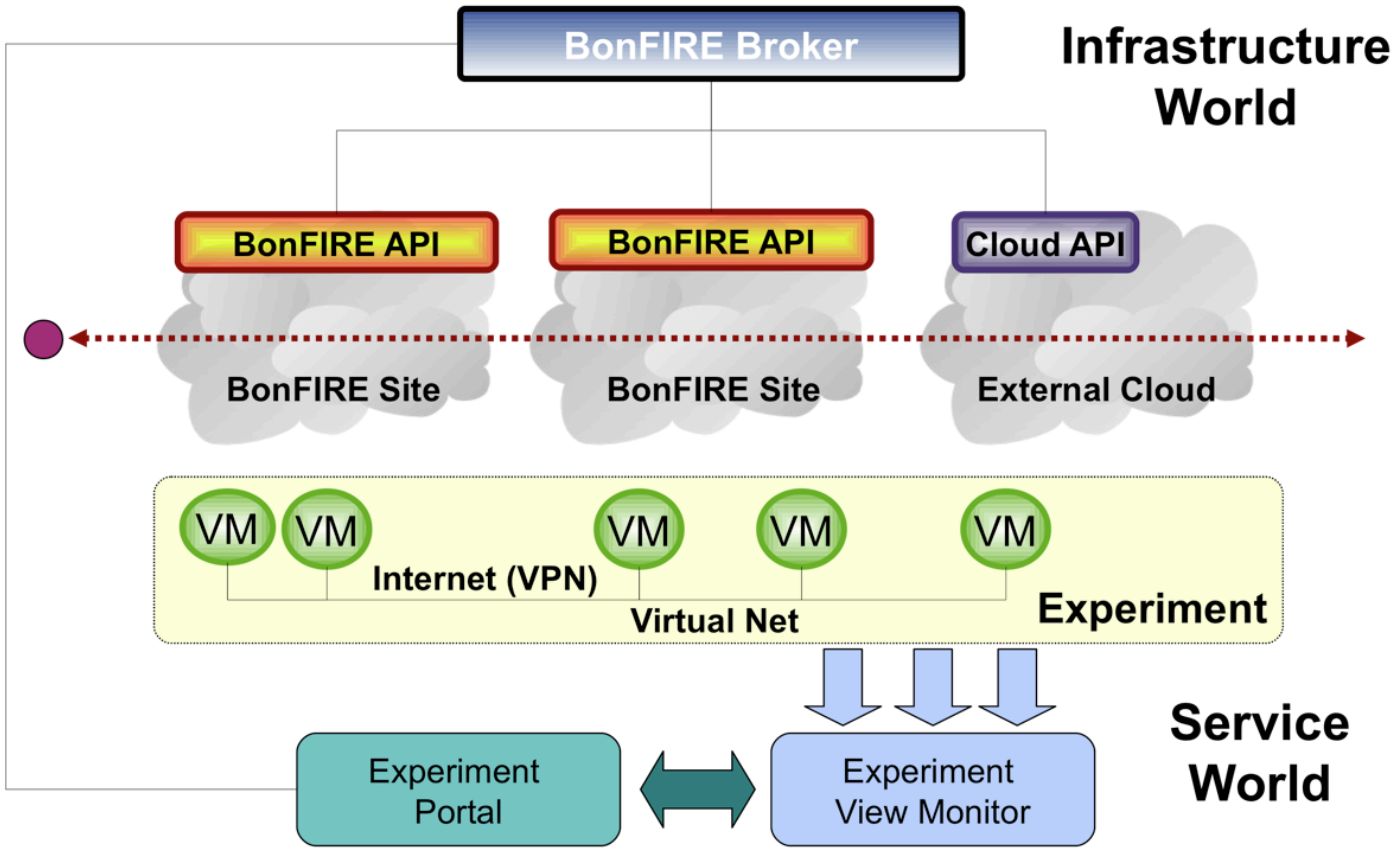

Source: BonFIRE Project

Tightly Coupled Federation – Multi-tier Architecture

Federation with a Cloud with Advanced Interoperation Support

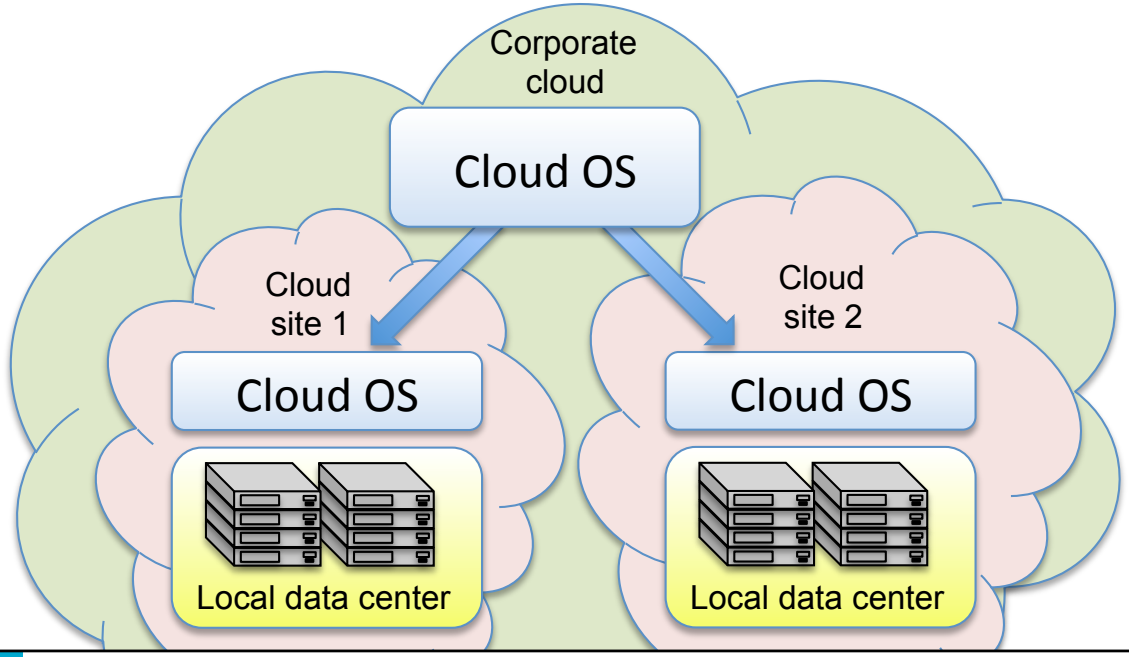

Aim	• Very large scale or geographically distributed data centers
Control	• Placement on specific physical resources • Same instance types
Monitoring & Accounting	• Physical resource consumption
Cross-site	• Live migration • High availability
Security	• User space sharing

Multi-tier Cloud Architecture

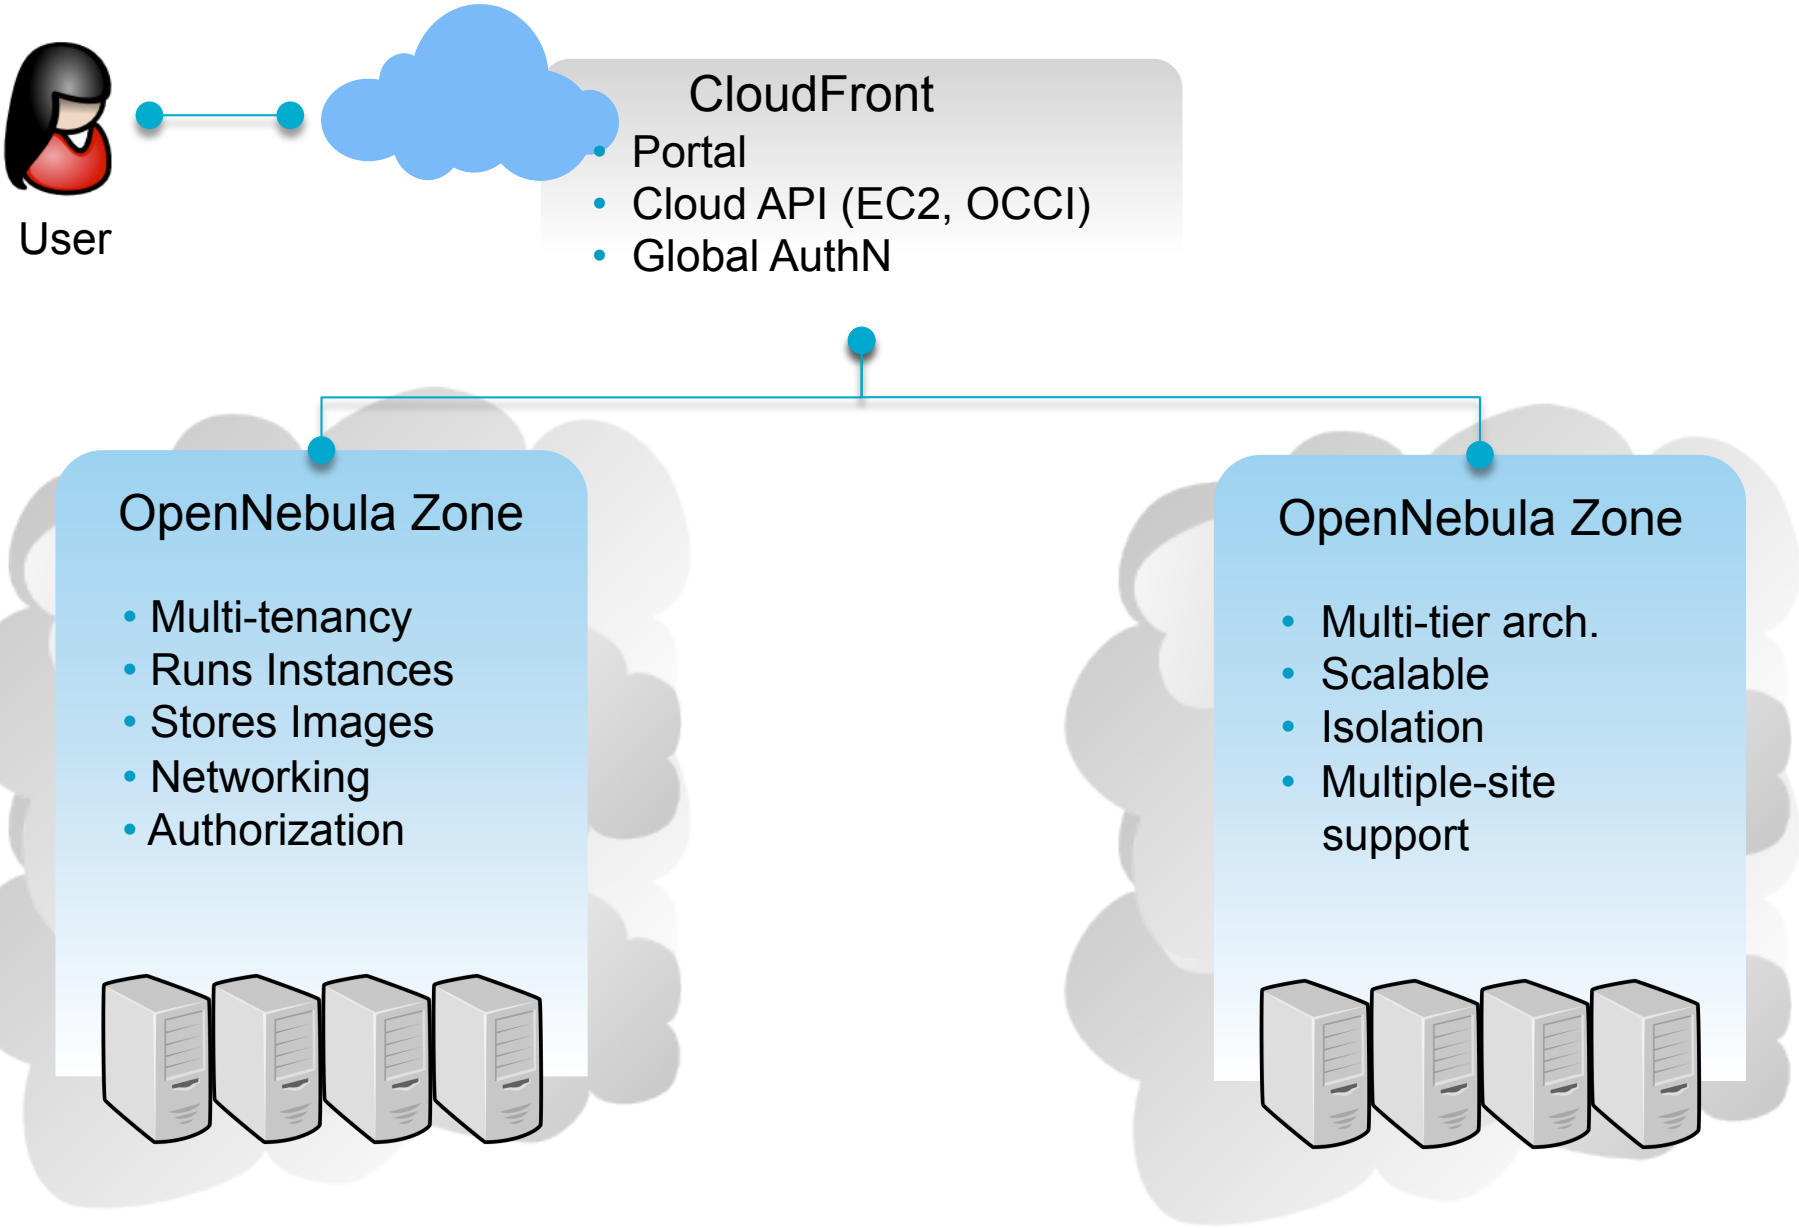

Challenges for Interoperability

Transparent Combination of Local Resources with Cloud Resources

1. Management Interfaces for Virtual Workloads

2. Management Interfaces for Data Elements

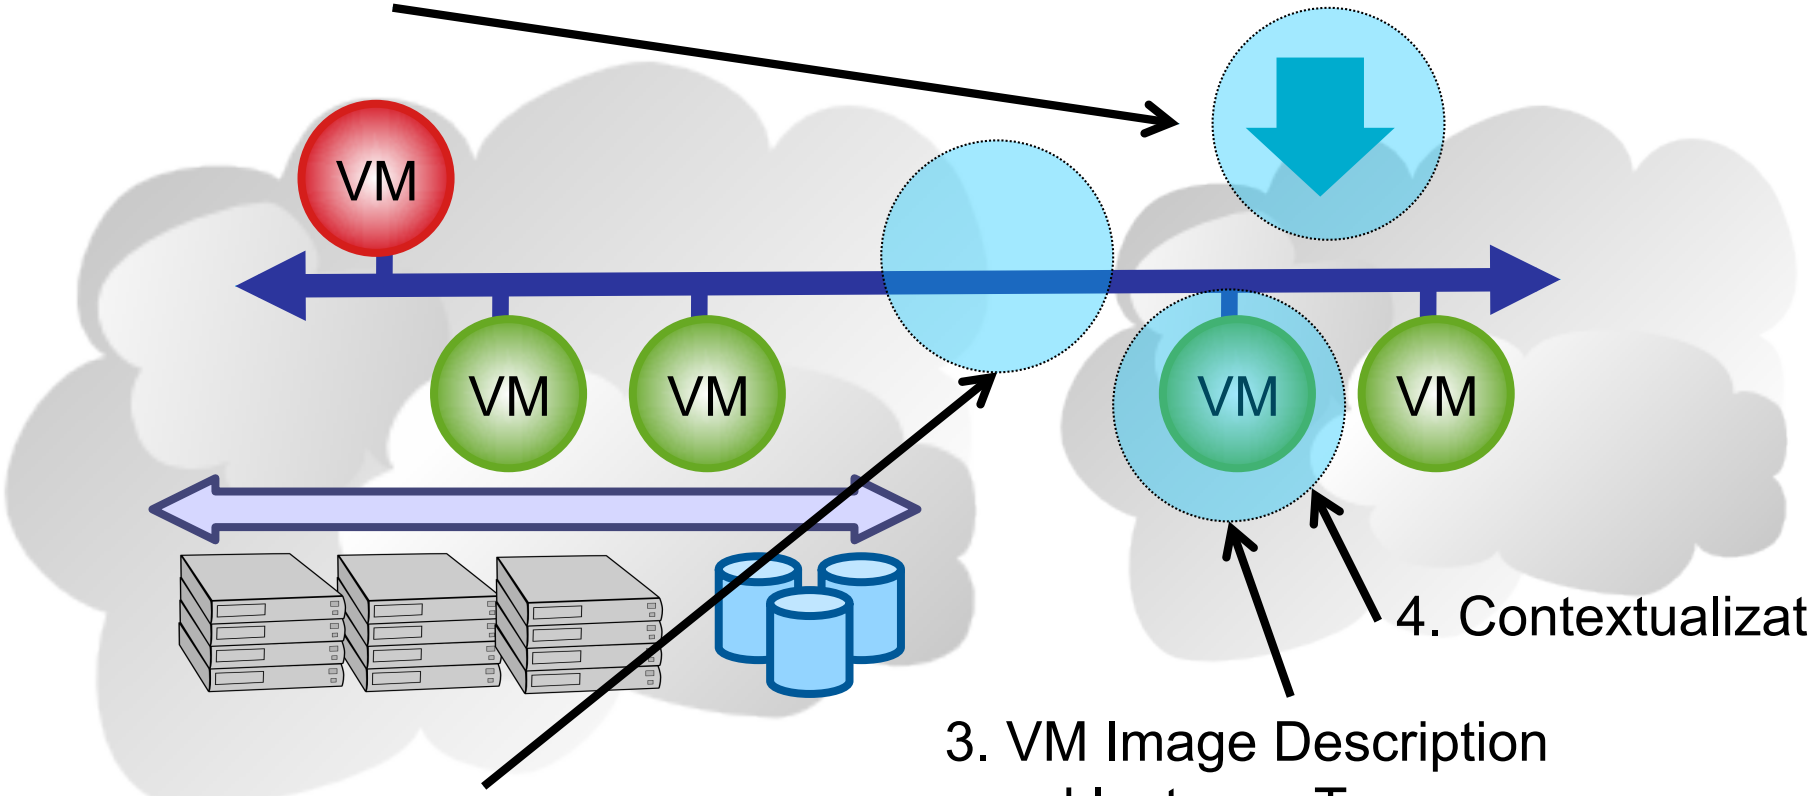

3. VM Image Description and Instance Type

4. Contextualization

5. Cross-site Networking and Latency

6. Common framework for security and QoS

Leveraging Existing Standards and Implementing Interoperation

Standardization

- Implement standards
- Integrate with standards

Which Standard?

- Different *de jure* standards
- Several *de facto* standards

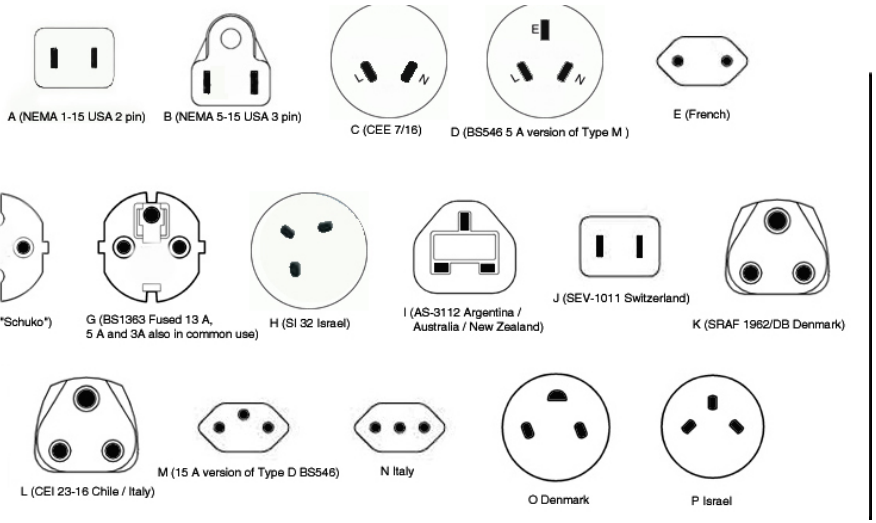

Interoperation

- Implement adaptors
- Use transformers

Grid and Cloud as Complementary Computing Models

Grids

Usage

- Job Processing
- Big Batch System
- File Sharing Services

Achievements

- Federation of Resources
- VO Concept

But…

- User experience
- Complexity

Clouds

Usage

- Raw infrastructure
- Elasticity & Pay-per-use
- Simple Web Interface

Achievements

- Agile Infrastructures
- IT is another Utility

But…

- Interoperability
- Federation

Resource Sharing

Scientific Applications

Uniform Security

Resource Management

Customize Environments

Flexibility & Simplicity

The OpenNebula Vision for Grid Sites

- Batch Job Processing
- Custom Execution Environments
- Grid Service Integration

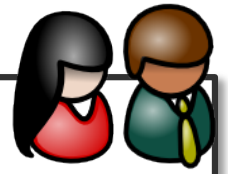

- Industry Applications
- Other WMS (pilots)
- Complete Services (cluster)

Acc

Grid Middleware

IaaS Interface

Service

LRMS (LSF, PBS…)

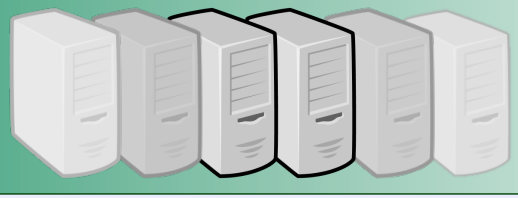

Virtual CE, WN…

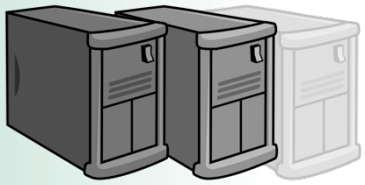

Other (web, mail...)

Raw machines

Provision

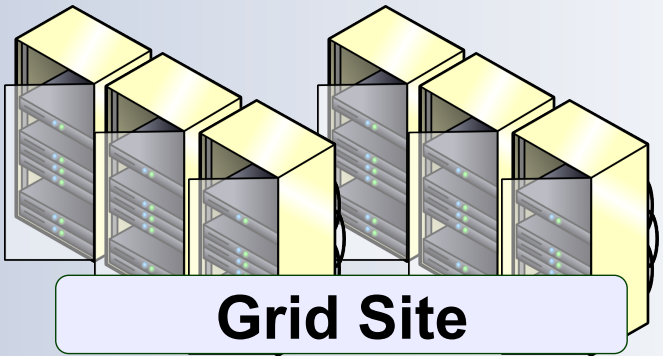

Grid Site

External Providers

The OpenNebula Vision for Grid Infrastructures

Grid Services

- Federation facilities
- Security
- Grid specific services

MarketPlace

- Sharing existing VM images
- Registry of metadata
- Image are kept elsewhere
- Supports trust

Appliance Repo

- Storage VM images
- Distributed
- Multi-protocol

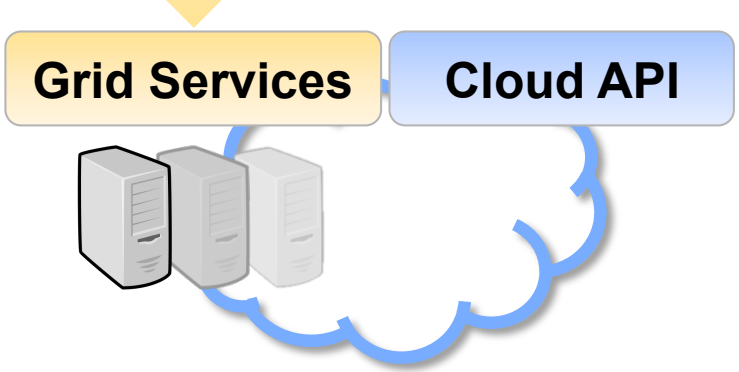

Cloud/Grid Site

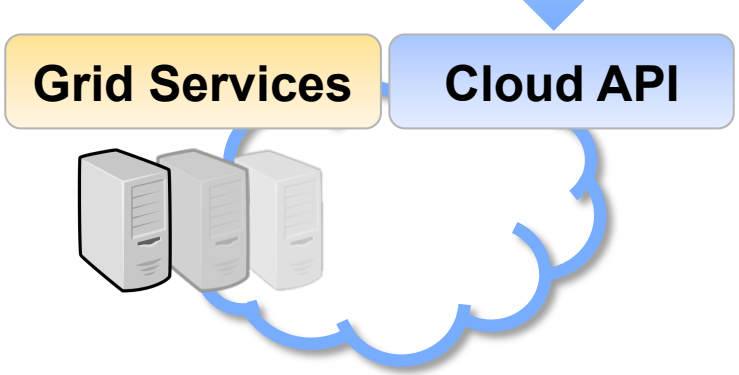

Cloud/Grid Site

OpenNebula is Driven by User Needs

How to contribute

- Join our mailing list
- Follow the development at dev.opennebula.org
- Contributions of by users: RIM, Akamai, Logica, FermiLab, SARA, Terradue, China Mobile… (approx. 100 listed at http://www.opennebula.org/about:contributors)
- 500 validated users at dev.opennebula

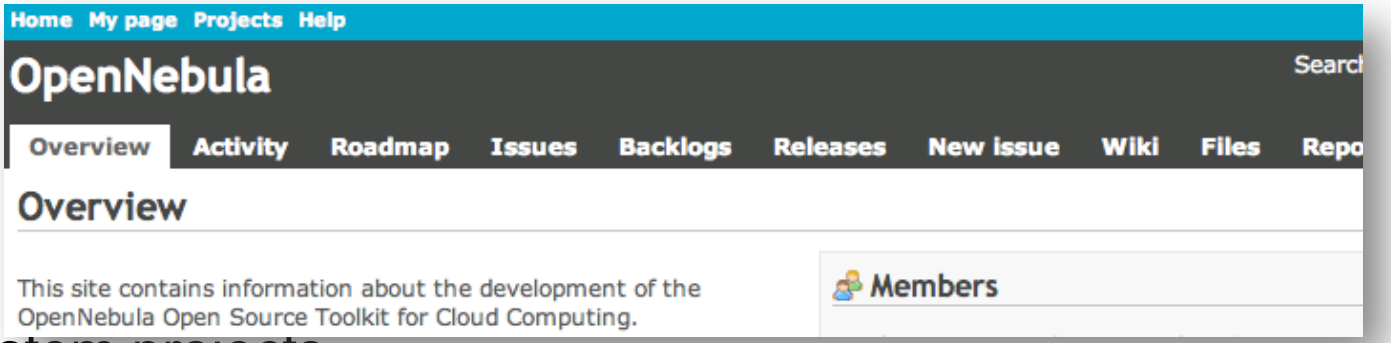

Ecosystem projects

- OpenNebula hosts an ecosystem catalog
- Promote and discuss ecosystem projects in our ecosystem mailing list

IRC Channel

- **#opennebula** on **irc.freenode.net**

Questions?

We Will Be Happy to Answer Any Question

CloudPlan.org

@imllorente

The research leading to these results has received funding from the *Ministerio de Ciencia e Innovación* of Spain through research grant TIN2009-07146.

OpenNebula vs. OpenStack?

Both Are Apache-licensed, Fully Open-source, Publicly Developed Technologies, but...

Feature	OpenStack	OpenNebula
Aim (Technical)	Public cloud (AWS-like deployments)	Private cloud & virtual datacenters (vCloud-like)
Functionality	Unique features for data center virtualization management, like VDCs, hybrid…	
Integration Capabilities	Very simple integration thanks to its plug-in based modular architecture by sys admins. OpenStack requires modifications in the code by experts	
Release Model	Developer community comprising different subprojects with different levels of maturity that require integration	Enterprise open-source product for computing clouds with a single installing, patching and updating process
Development Model	Consensus-based approach where vendors try to meet the needs of the project and their monetization goals	User-driven development with contributions from the users