

**Grids & e-Science 2009
UIMP 2009 Summer School
June 16th, 2009**

An Introduction to Cloud Computing Infrastructures

Ignacio M. Llorente

dsa-research.org

**Distributed Systems Architecture Research Group
Universidad Complutense de Madrid**



Contents

An Introduction to Cloud Computing Infrastructures

- Cloud computing as **resource provisioning platform**
- **Types** of cloud infrastructures
- **Middleware** for cloud computing
- **Multi-cloud** architectures
- Building cloud infrastructures with **OpenNebula**
- **Research** in cloud computing



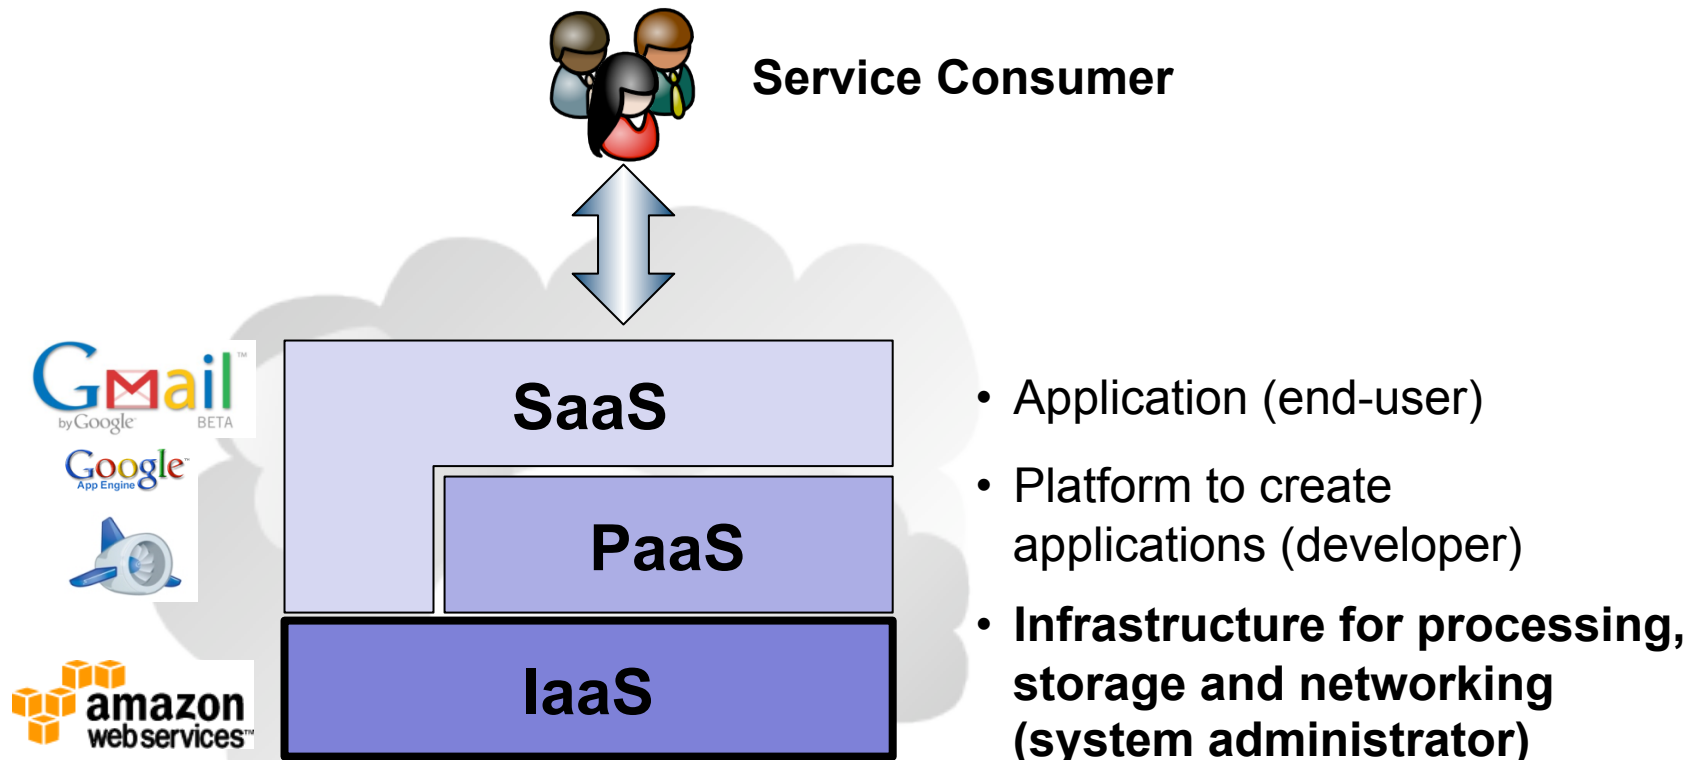
Cloud Computing as a Resource Provisioning Platform

An Introduction to Cloud Computing Infrastructures

What is a Cloud Computing Infrastructure?

- Resource provisioning platform offering on-demand and ubiquitous network access to elastic and configurable IT-enabled capabilities delivered as a service

Delivery Models in Cloud Computing

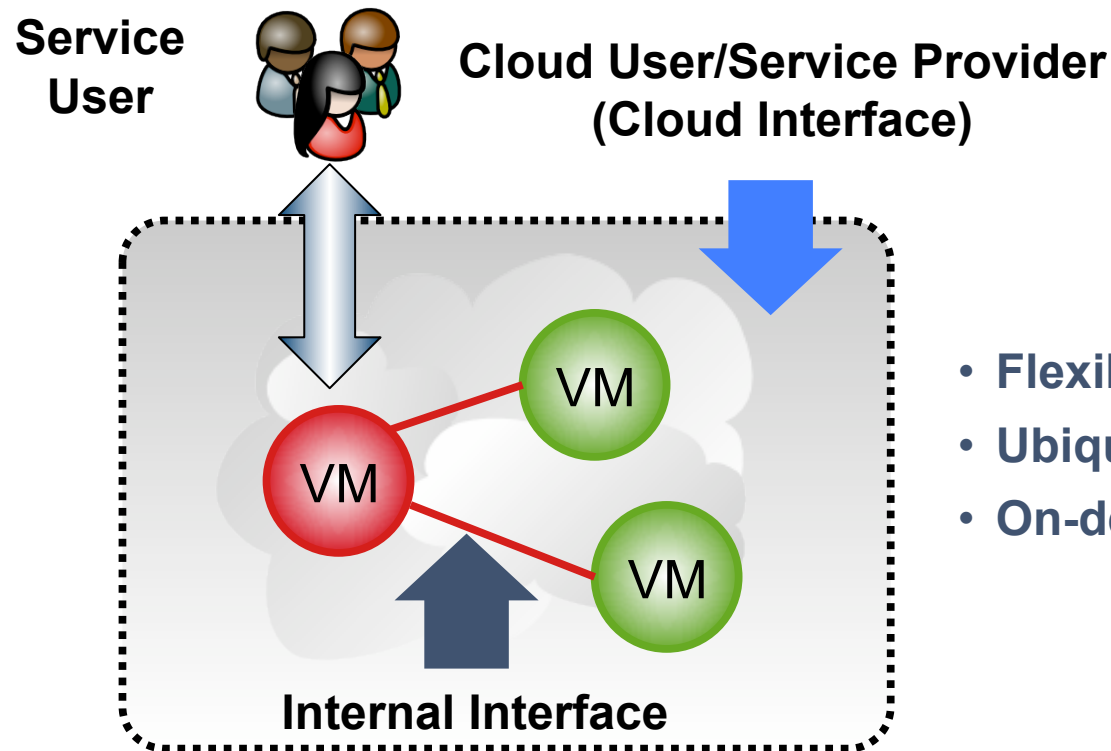


Cloud Computing as a Resource Provisioning Platform

An Introduction to Cloud Computing Infrastructures

Infrastructure as a Service

- Platform for execution of **virtualized services (groups of interconnected VMs)**
- **Internal interface:** Full management of virtual and physical resources
- **Cloud interface:** Simple and remote management of cloud (virtual) resources at a higher abstraction level



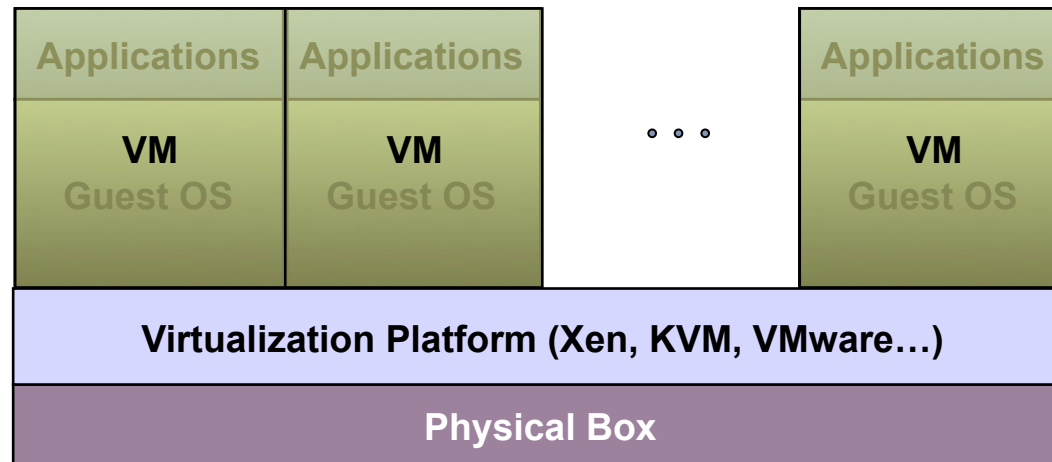
- **Flexible and elastic capacity**
- **Ubiquitous network access**
- **On-demand access**

Cloud Computing as a Resource Provisioning Platform

An Introduction to Cloud Computing Infrastructures

Virtualization is the Foundation of Infrastructure Clouds

- A VM is an **isolated runtime environment** (guest OS and applications)
- **Multiple virtual systems (VMs)** to run on a single physical system
- **Separation** of service/server from the physical resources



Types of Cloud Infrastructures

An Introduction to Cloud Computing Infrastructures

Commercial Public Clouds

- The infrastructure is **owned by a single commercial organization and used by customers**
- Public clouds enable the deployment of an entire IT infrastructure **without the associated capital costs, paying only for the used capacity**
- **Cloud interface:** Simple remote management of virtualized server instances

Examples of Commercial Providers

Commercial Cloud Tools



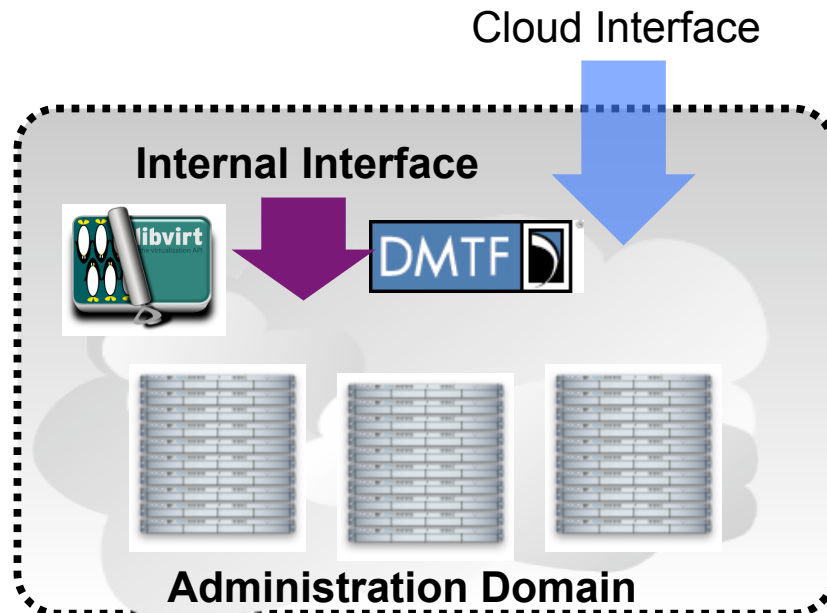
Types of Cloud Infrastructures

An Introduction to Cloud Computing Infrastructures

Private Clouds

- **The infrastructure is owned and used by a single organization**
- **Private clouds enable** a flexible and agile management of local infrastructure
- **Not a new model**, datacenter management has been around for a while
- **Internal interfaces expose additional functionality** for managing virtualized resources and controlling data center operation, not exposed by cloud interfaces
- **Cloud interfaces** may be also provided for users requiring higher abstraction

Benefits



- **Centralized management**
- **VM placement optimization**
- **Dynamic resizing and partitioning of the infrastructure**
- **Support for heterogeneous workloads**

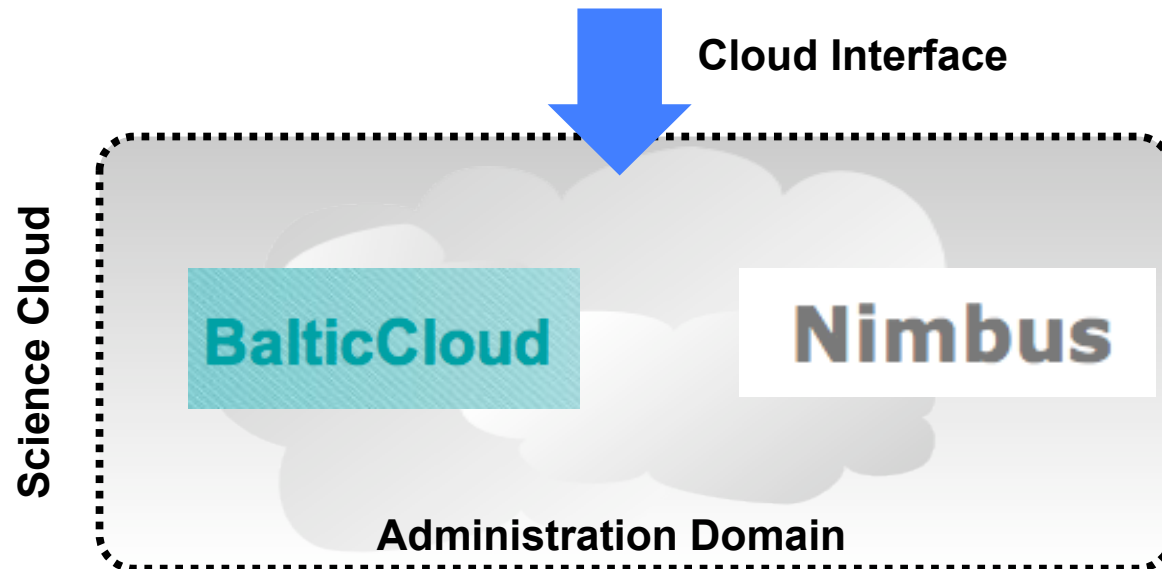
Types of Cloud Infrastructures

An Introduction to Cloud Computing Infrastructures

Domain-specific Clouds

- **Scientific Clouds** are clouds designed to offer specific services for the scientific community, usually to enable scientific and educational projects to experiment with cloud computing
- **HPC Clouds** are specifically designed for tightly-coupled applications

Examples of Scientific Clouds



Middleware for Cloud Computing

An Introduction to Cloud Computing Infrastructures

Open-Source Technologies



Eucalyptus

- Compatible with **Amazon EC2 interfaces** and designed to support additional client-side interfaces



the globus[®] alliance

Nimbus

- **Exposes EC2 and WSRF interfaces and offers self-configuring virtual cluster support**



enomaly
elastic computing

- Management of clusters and virtual machines through a single **easy to use interface**

OpenNebula

- **Scalable, flexible and dynamic management** of virtual infrastructures within and across sites to build any cloud computing deployment model

Commercial Software

VMware vSphere™

- **VMware** solution for private cloud computing

Citrix Cloud Center

- **Citrix** solution for private cloud computing

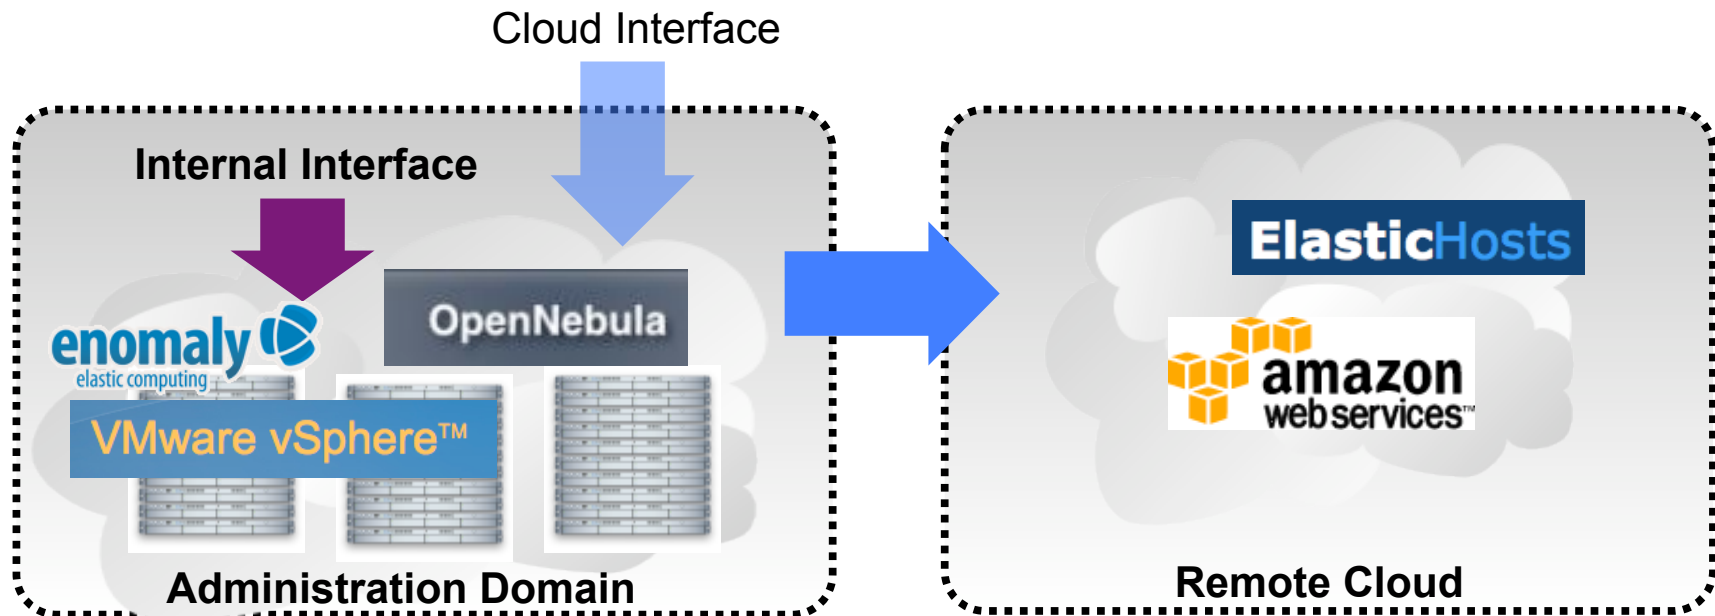
Multi-Cloud Architectures

An Introduction to Cloud Computing Infrastructures

Hybrid Clouds

- **Extension of Private Clouds** to combine private with public Cloud-based infrastructure to enable highly scalable hosting environments
- Access to remote cloud is **fully transparent** to private cloud users
- Hybrid Clouds enable the **dynamic scaling of capacity to meet peak or fluctuating demands**

Infrastructure Elasticity




Multi-Cloud Architectures

An Introduction to Cloud Computing Infrastructures

Federation of Clouds

- **Extension of Hybrid Clouds** to support sharing of resources between partners
- Access to remote clouds is **fully transparent** to cloud users
- **No single facility/provider can create a seemingly infinite infrastructure** capable of serving massive amounts of users at all times, from all locations
- **Different levels of inter-site** coupling in terms of contract type, information disclosure, and degree control over remote hosts

Interoperability Across Infrastructures



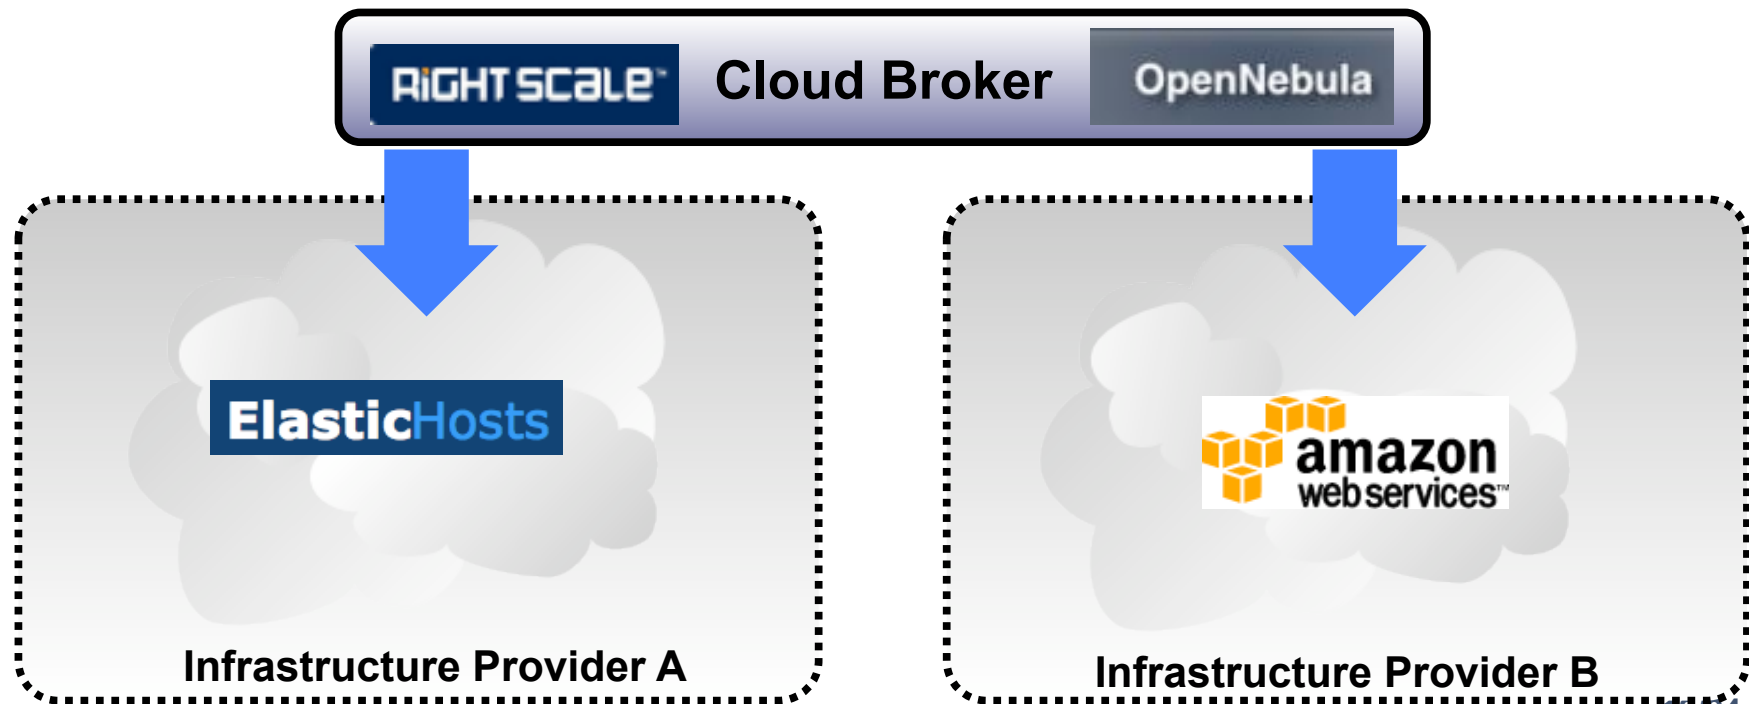
Multi-Cloud Architectures

An Introduction to Cloud Computing Infrastructures

Cloud Brokering

- **Simultaneous use of different cloud providers**
- Provisioning of resources from different clouds to improve the **cost-effectiveness** of the deployment, to **reduce the response time for user requests**, or to implement **high-availability** strategies

Multi-Cloud Infrastructures

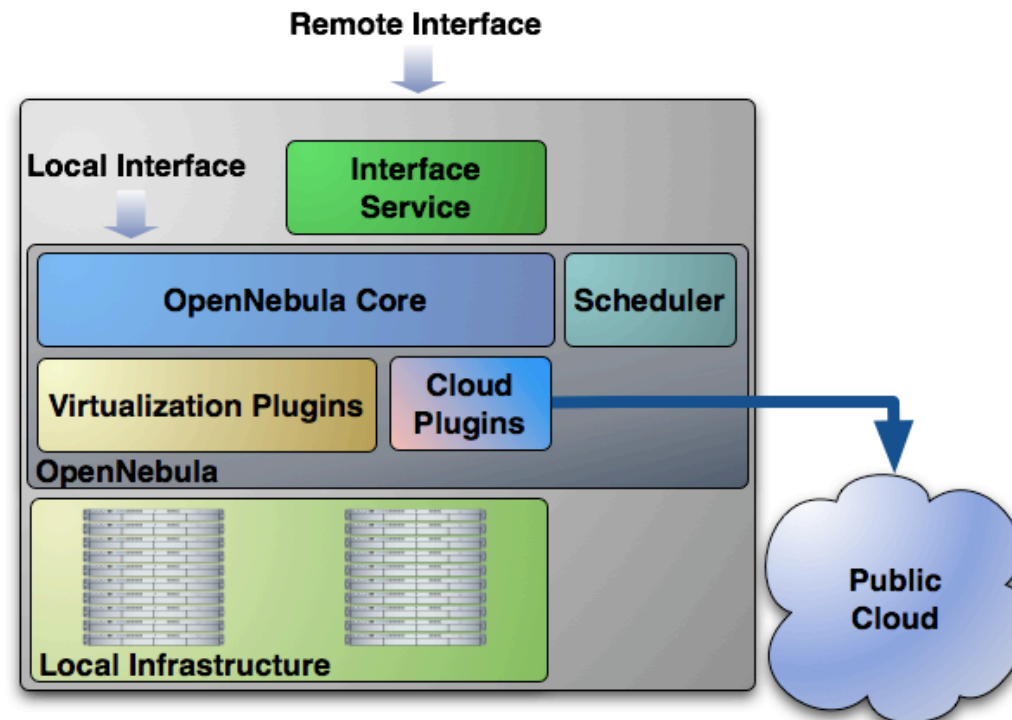


Building Cloud Infrastructures with OpenNebula

An Introduction to Cloud Computing Infrastructures

The Open-Source Toolkit for Building Cloud Infrastructures

- **Dynamic deployment and re-placement of virtualized services** (groups of interconnected virtual machines) within and across sites
- Its components can be used **to build any type of cloud deployment**: private, hybrid, and public clouds; and cloud brokers





Building Cloud Infrastructures with OpenNebula

An Introduction to Cloud Computing Infrastructures

Features

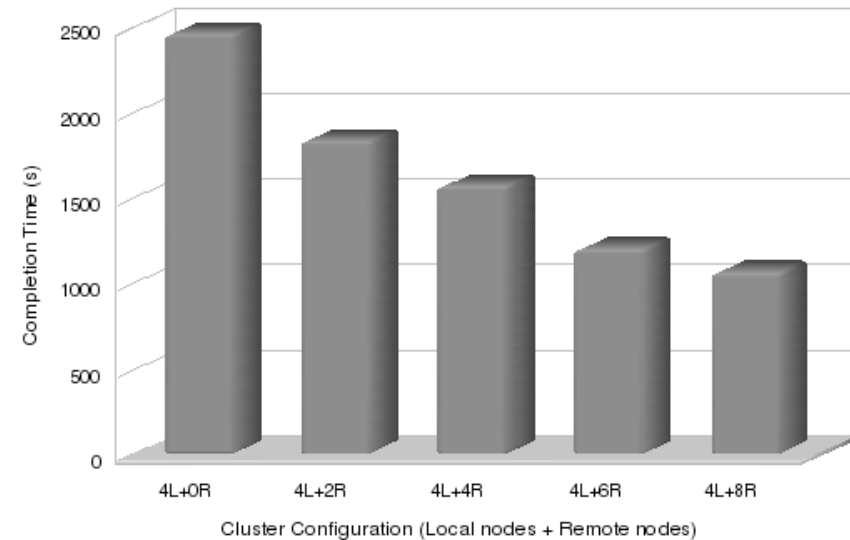

| Feature | Function |
|----------------------------------|--|
| Internal Interface | <ul style="list-style-type: none">• Unix-like CLI for fully management of VM life-cycle and physical boxes• XML-RPC API and libvirt interface |
| Cloud Interface | <ul style="list-style-type: none">• RESTful service for basic management of virtualized resources |
| Scheduler | <ul style="list-style-type: none">• Requirement/rank matchmaker• Generic framework to build any scheduler for load balancing and energy efficient server consolidation |
| Virtualization Management | <ul style="list-style-type: none">• Xen, KVM, VirtualBox, VMware and libvirt connectors• Amazon EC2 |
| Image Management | <ul style="list-style-type: none">• General mechanisms to transfer and clone VM images |
| Network Management | <ul style="list-style-type: none">• Definition of isolated virtual networks to interconnect VMs |
| Fault Tolerance | <ul style="list-style-type: none">• Persistent database backend to store host and VM information |
| Contextualization | <ul style="list-style-type: none">• Support for automatic configuration of service components |
| Scalability | <ul style="list-style-type: none">• Tested in the management of hundreds of VMs |
| Installation | <ul style="list-style-type: none">• Installation on a UNIX cluster front-end without requiring new services in the remote resources• Distributed in Ubuntu 9.04 (Jaunty Jackalope), due in April 2009 |
| Flexibility | <ul style="list-style-type: none">• Open and flexible architecture and interfaces, open source software |

Building Cloud Infrastructures with OpenNebula

An Introduction to Cloud Computing Infrastructures

Example: Computing Cluster on Hybrid Cloud

- Elastic execution of a **SGE** computing cluster
- **Dynamic scaling** of the number of worker nodes to meet demands
- **Private network** with NIS and NFS; and VPN across sites


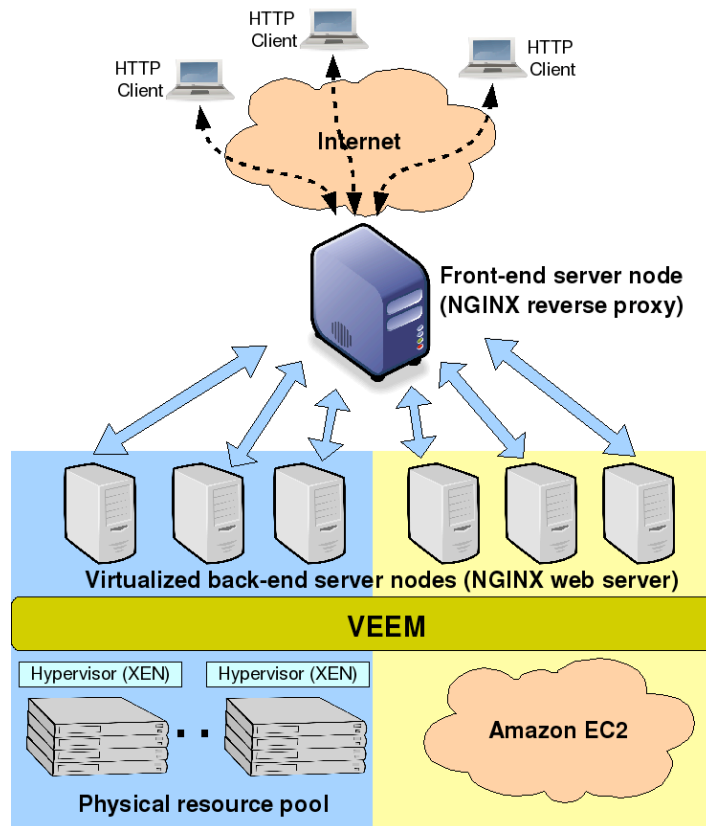


Building Cloud Infrastructures with OpenNebula

An Introduction to Cloud Computing Infrastructures

Example: Web Server on Hybrid Cloud

- 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


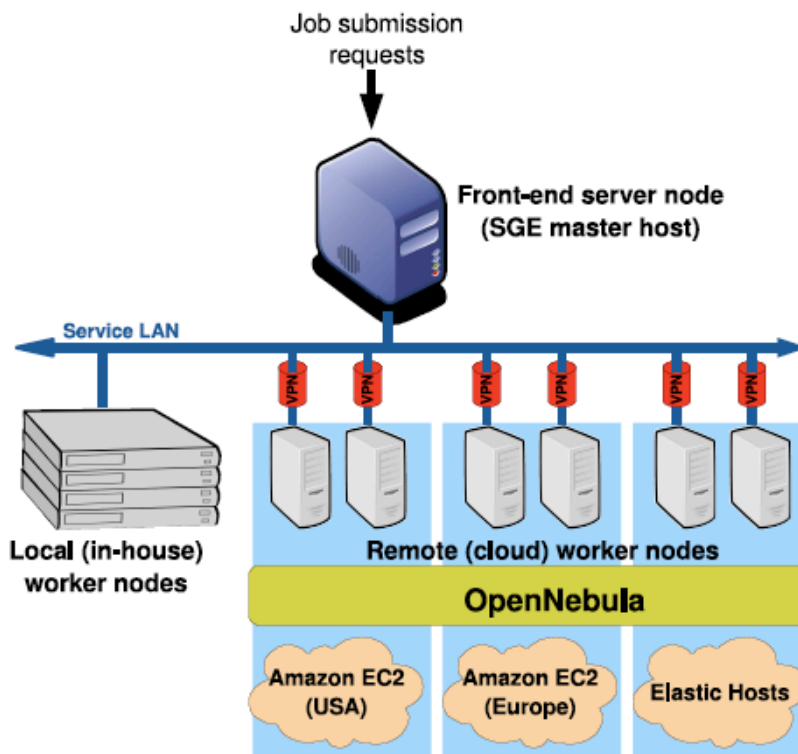


Building Cloud Infrastructures with OpenNebula

An Introduction to Cloud Computing Infrastructures

Example: Cloud Brokering of Computing Cluster

- Access to three different clouds to dynamically optimize the performance/cost ratio of a computing cluster
- The optimal deployment depends on the virtualized service profile and even on its workload



Building Cloud Infrastructures with OpenNebula

An Introduction to Cloud Computing Infrastructures

The Ecosystem

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

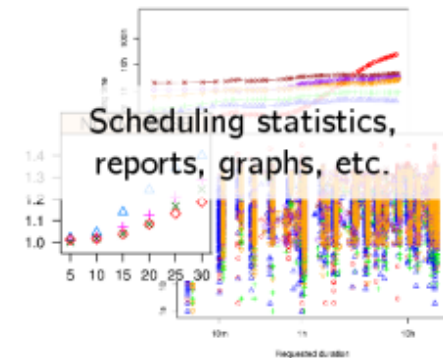
 **HAIZEA**

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



With
OpenNebula

Simulation



- **Experimental VM Consolidation Scheduler for Energy Efficiency :** Periodic re-placement of VMs for server consolidation and suspension/resume of physical resources
- **Nimbus:** Can be used as a WSRF or EC2 front-end.
- **ElasticHosts Plug-in:** Meet fluctuating peak demands using this cloud provider

Research in Cloud Computing

An Introduction to Cloud Computing Infrastructures

Main Research Projects

Open Cirrus

- **Research testbed** designed to support research into the design, provisioning, and management of services at a global, multi-datacenter scale

RESERVOIR

- Investigate **technologies** for advanced Cloud Computing

StratusLab

- Investigate the application of cloud computing to **Grid infrastructures**

Cloud Computing Brings New Research Challenges

- **Research in Cloud Computing addresses the new challenges** of:
 - Meeting the **technology requirements** of next generation private, public and hybrid cloud computing architectures
 - Allowing **service workloads** to take advantage of cloud computing
- Because cloud solution architectures include technology **components from different fields**, many research challenges in Cloud Computing have been already addressed to a certain degree by **different research communities**

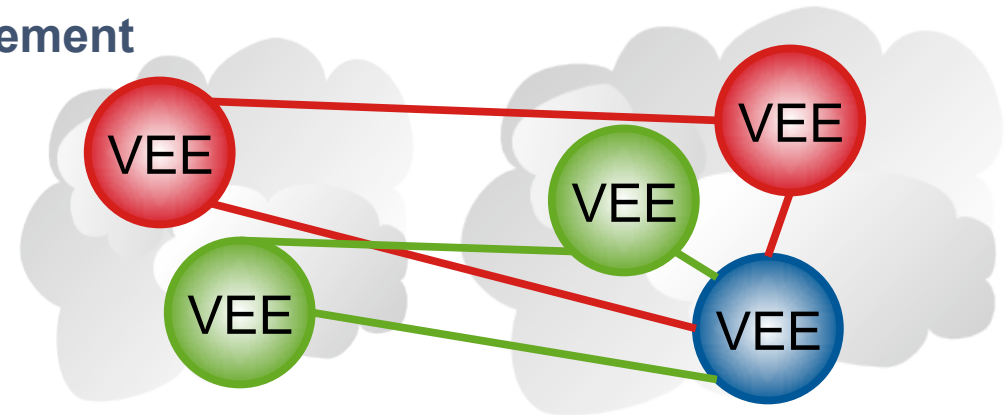
Research in Cloud Computing

An Introduction to Cloud Computing Infrastructures



Resources and Services Virtualization without Barriers

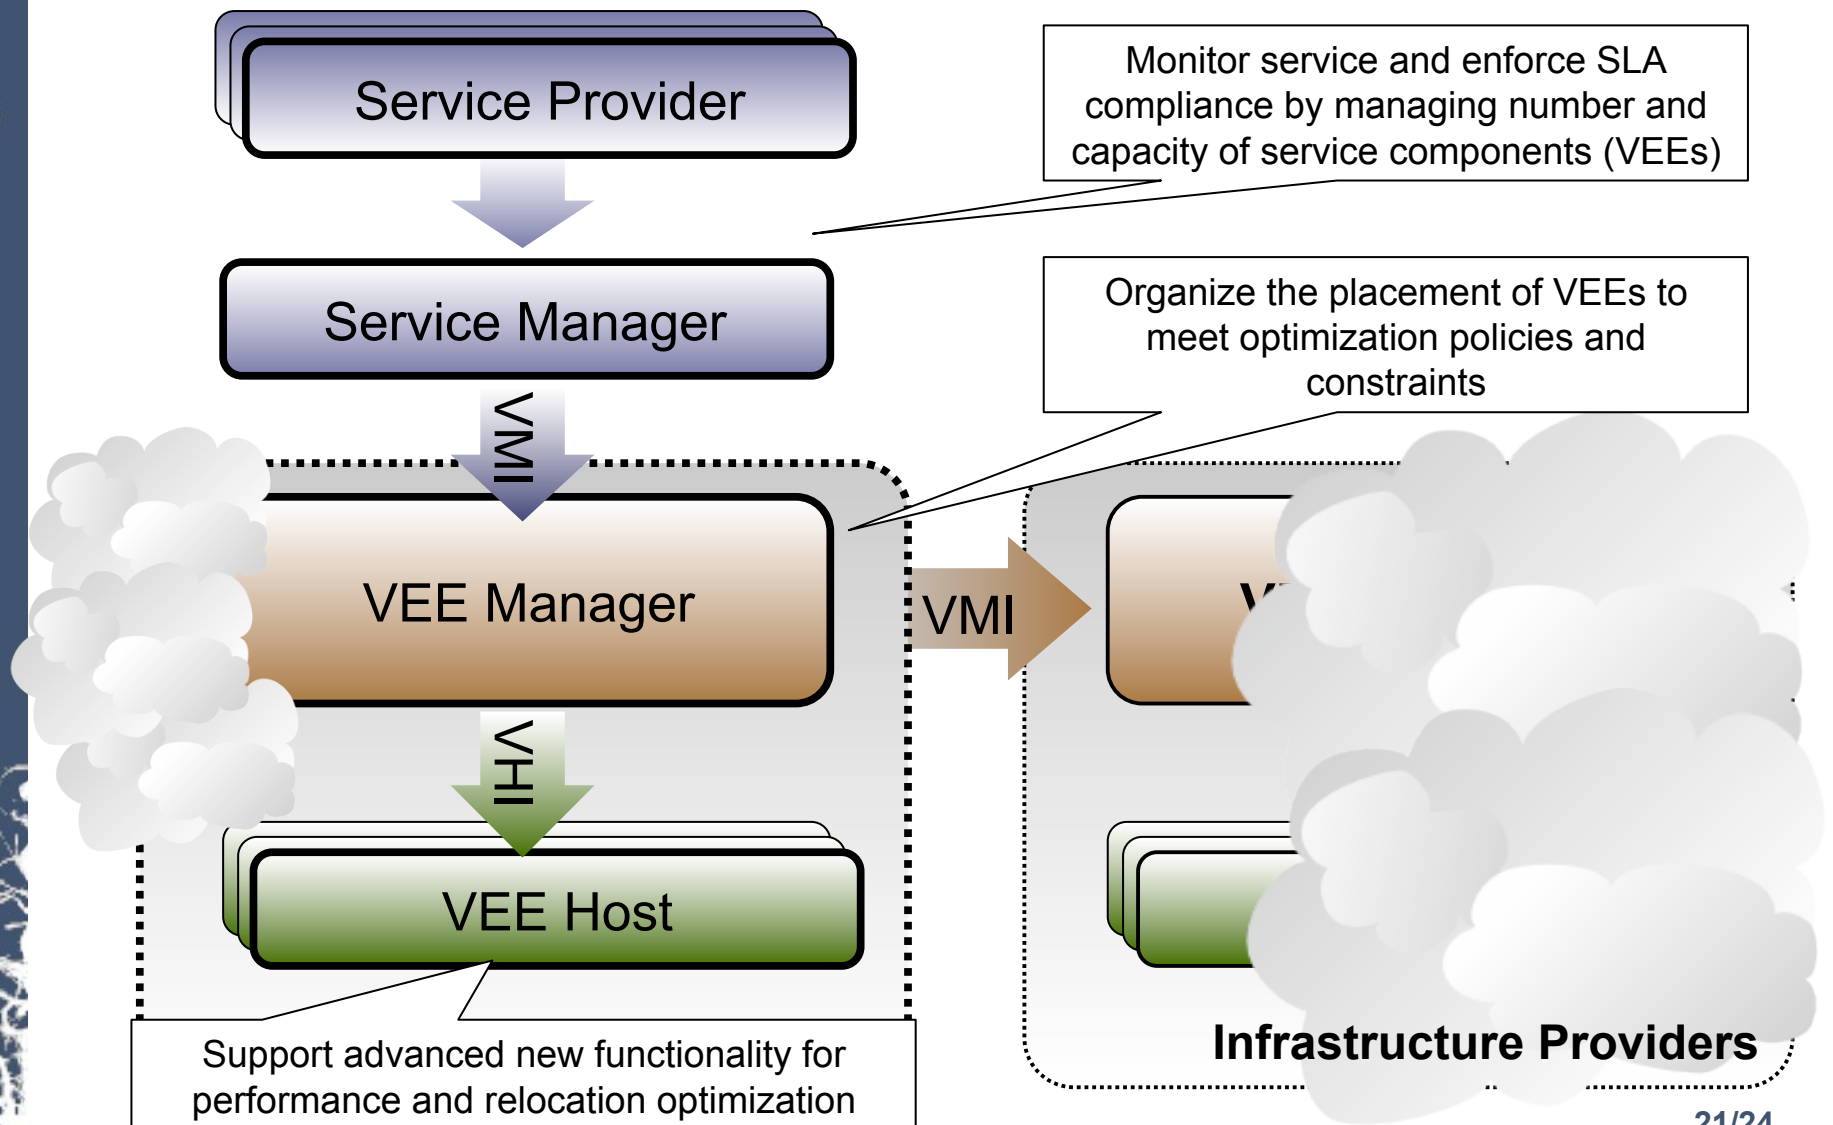
- **Open source technology** to enable deployment and management of **complex IT services across different administrative domains**
- **Functionality for Service Management**
 - Definition
 - Lifecycle
 - Billing/accounting
 - Elasticity/SLAs



Flagship of Cloud Computing Research in FP7

- Focus on technologies that enable to build a **federation of cooperating computing clouds**
- **A project driven by business use cases:** SAP business application, Telco application, utility computing and eGov application
- 17-million and 3-year project partially **funded by the European Commission** (NESSI Strategic Project)

RESERVOIR Architecture



Research in Cloud Computing

An Introduction to Cloud Computing Infrastructures




Enhancing Grid Infrastructures with Cloud Technologies

- This project focuses on the **benefits from using cloud computing on Grid** by:
 - Exploring **the integration of cloud technologies** and services, especially virtualization, into existing Grid Infrastructures.
 - Investigating how existing Grid infrastructures can provide a **cloud-like delivery model** as well as existing Grid services

Grid in the Cloud

- Technical, operational and security challenges raised in the **deployment of a Grid site in the Cloud**





My Vision on the Future of Cloud Computing

An Introduction to Cloud Computing Infrastructures

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
- **Public Cloud providers will implement 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



An Introduction to Cloud Computing Infrastructures

You can create now your cloud infrastructure

OpenNebula is available for download

More info, downloads, mailing lists at
www.OpenNebula.org



www.reservoir-fp7.eu/

dsa-research.org

Acknowledgements

- Rubén S. Montero
- Tino Vazquez
- Javier Fontan
- Rafael Moreno
- Raúl Sampedro