New Paradigms: Clouds, Virtualization and Co.
EGEE08, Istanbul, September 25, 2008

An Introduction to Virtualization and Cloud Technologies to Support Grid Computing

Ignacio M. Llorente

dsa-research.org

Distributed Systems Architecture Research Group
Universidad Complutense de Madrid

Objectives

- **Introduce virtualization and cloud** from the perspective of the Grid community
- Show the **benefits of virtualization and cloud** for Grid computing
- Demonstrate how Grid, virtualization and cloud are complementary technologies that **will cooperate in future Grid computing infrastructures**
- Introduce the **RESERVOIR project**, European initiative in virtualization and cloud computing

Barriers for Adoption of the Compute Grid Model

- **High degree of heterogeneity** (software & hardware)
- **High operational costs**
- **Isolate and partition amount of resources** contributed to the Grid
- **Specific environment** requirements for different VOs

Grids are difficult to **mantain, operate and use**

Virtualization Platform

Separation of Virtual Machine from Physical Infrastructure

- **A VM is an isolated runtime environment** (guest OS and applications)
- **Multiple virtual systems (VMs)** to run on a single physical system

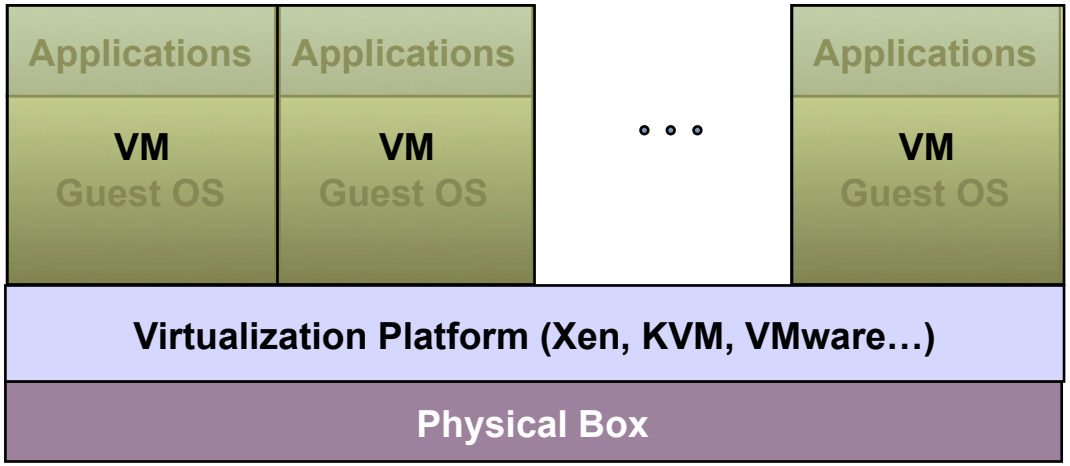

Benefits of Virtualization Platforms

- Natural way to deal with the **heterogeneity** of the infrastructure
- Allow **partitioning and isolating** of physical resources
- Execution of **legacy applications**

Distributed Management of VMs

An Introduction to Virtualization and Cloud Technologies to Support Grid Computing

Extending the Benefits of Virtualization to a Physical Cluster

- VM Managers creates a **distributed virtualization layer**
  - Extend the benefits of VM Monitors from one to multiple resources
  - Decouple the VM (service) from the physical location
- Transform a distributed physical infrastructure into a **flexible and elastic virtual infrastructure**

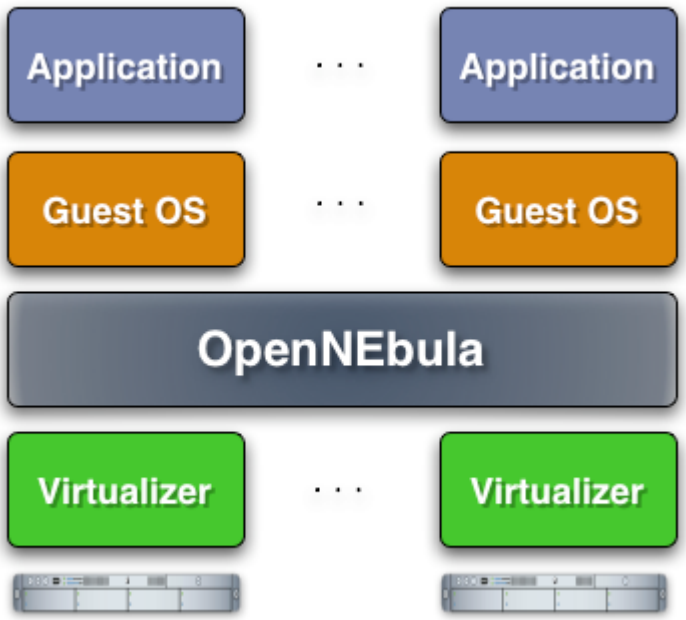

Benefits of VM Managers

- Centralized management
- Balance of workload
- Server consolidation
- Dynamic resizing of the infrastructure
- Dynamic cluster partitioning
- Support for heterogeneous workloads
- On-demand provision of VMs

Virtualization of a Computing Cluster

Separation of Resource Provisioning from Job Management

- New virtualization layer **between the service and the infrastructure layers**
- **Seamless integration** with the existing middleware stacks.
- **Completely transparent** to the computing service and so end users

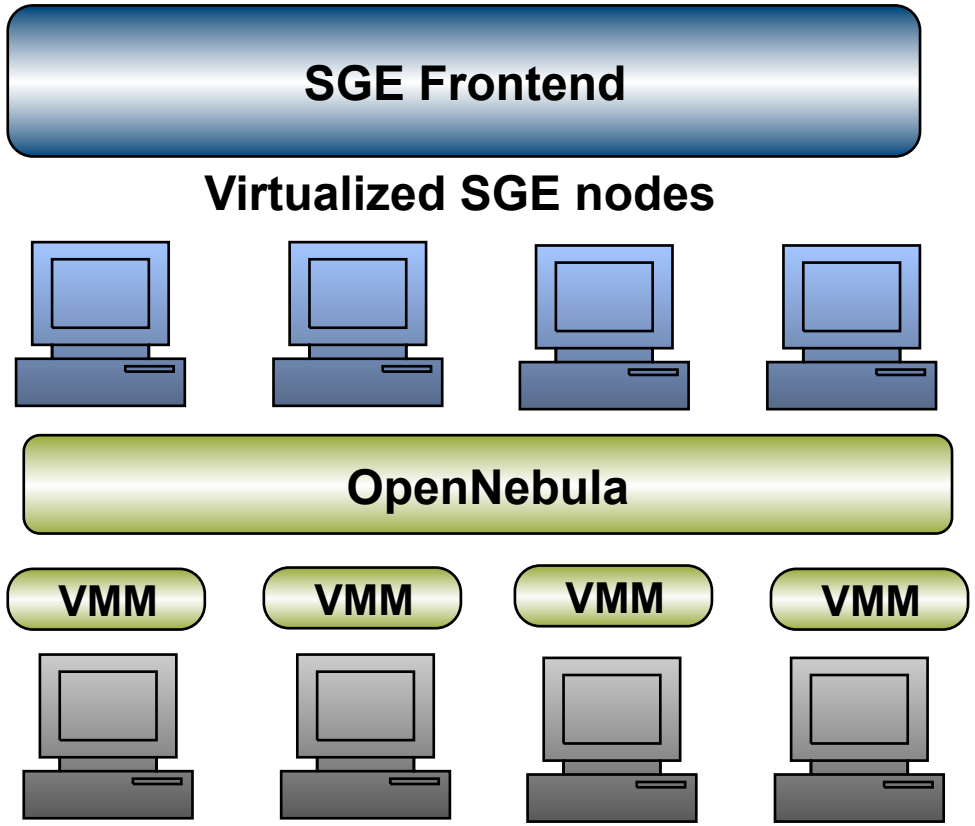

Integration of a Virtualized Cluster within a Grid

An Introduction to Virtualization and Cloud Technologies to Support Grid Computing

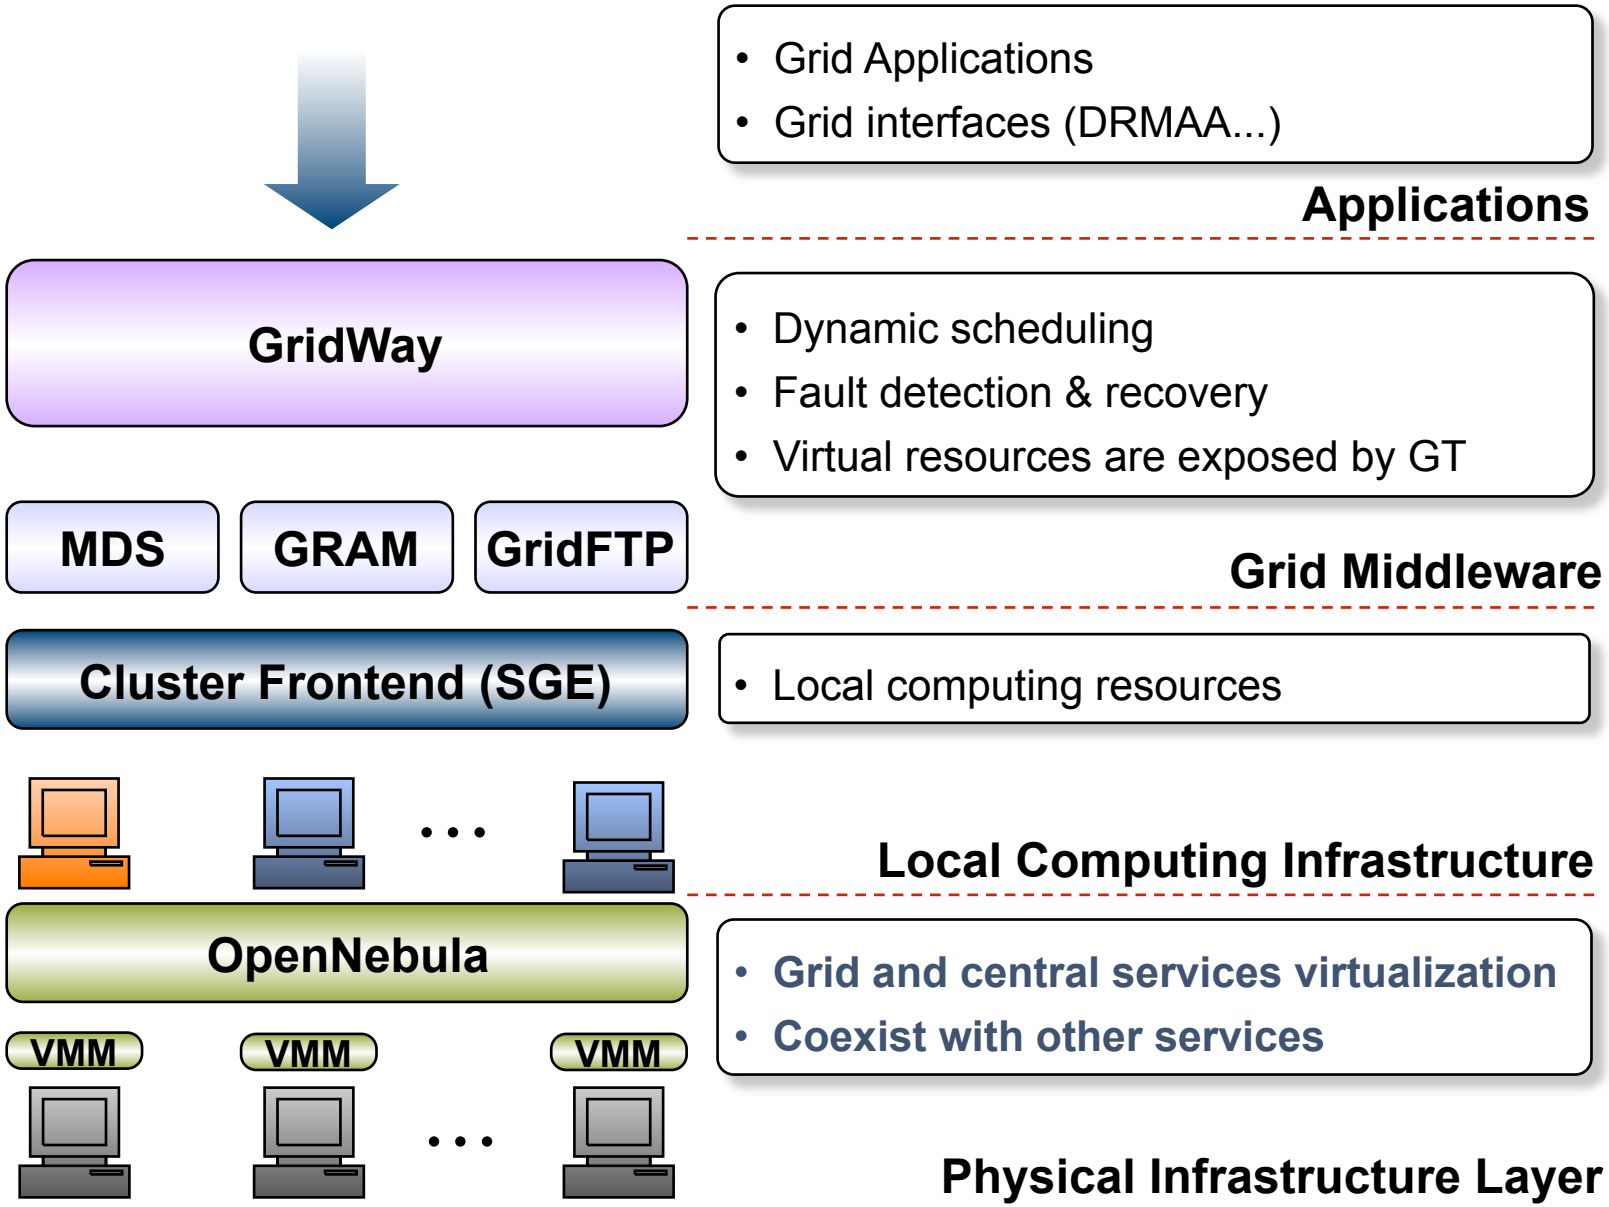

Integration of a Virtualized Cluster within a Grid

Benefits of Virtualization for Existing Grid Infrastructures

- The **virtualization of the local infrastructure** provides:
  - Easy support for VO-specific worker nodes
  - Reduce *gridification* cycles
  - Dynamic balance of resources between VO's
  - 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
  - Performance partitioning between local and grid services

Solve many of the obstacles for Grid adoption

Cloud as Provision of Virtualized Resources

A Service to Provide Hardware on Demand (IaaS)

- Cloud systems provide **virtualized resources as a service**
- Provide **remote on-demand access to infrastructure** for the execution of virtual machines

Simple Interfaces for VM Management

- Submission
- Control
- Monitoring

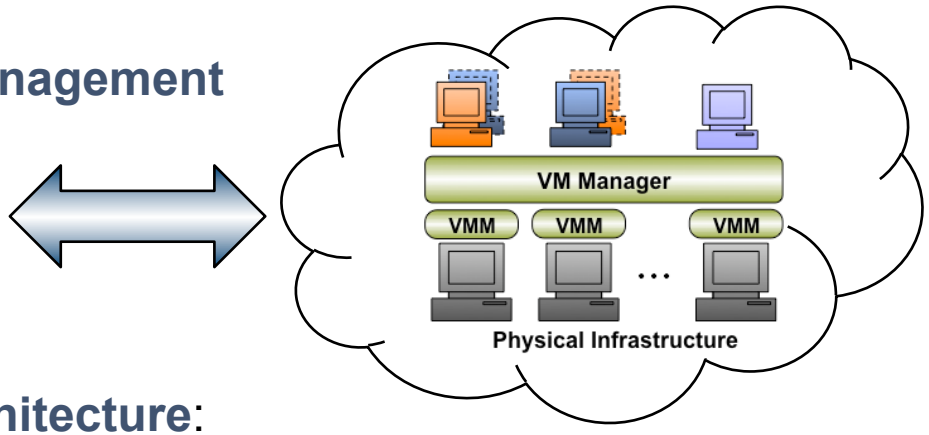

- Main components of a **Cloud architecture**:
  - Front-end: Remote interface (Eucalyptus, Globus Nimbus…)
  - Back-end: Local VM manager (OpenNebula)

Infrastructure Cloud Services

- **Commercial Cloud**: Amazon EC2, GoGrid, Flexiscale…
- **Scientific Cloud**: Nimbus (University of Chicago)

Cloud for Scaling out Local Infrastructures

An Introduction to Virtualization and Cloud Technologies to Support Grid Computing

On-demand Access to Cloud Resources

- Supplement local resources with cloud resources to **satisfy peak or fluctuating demands**

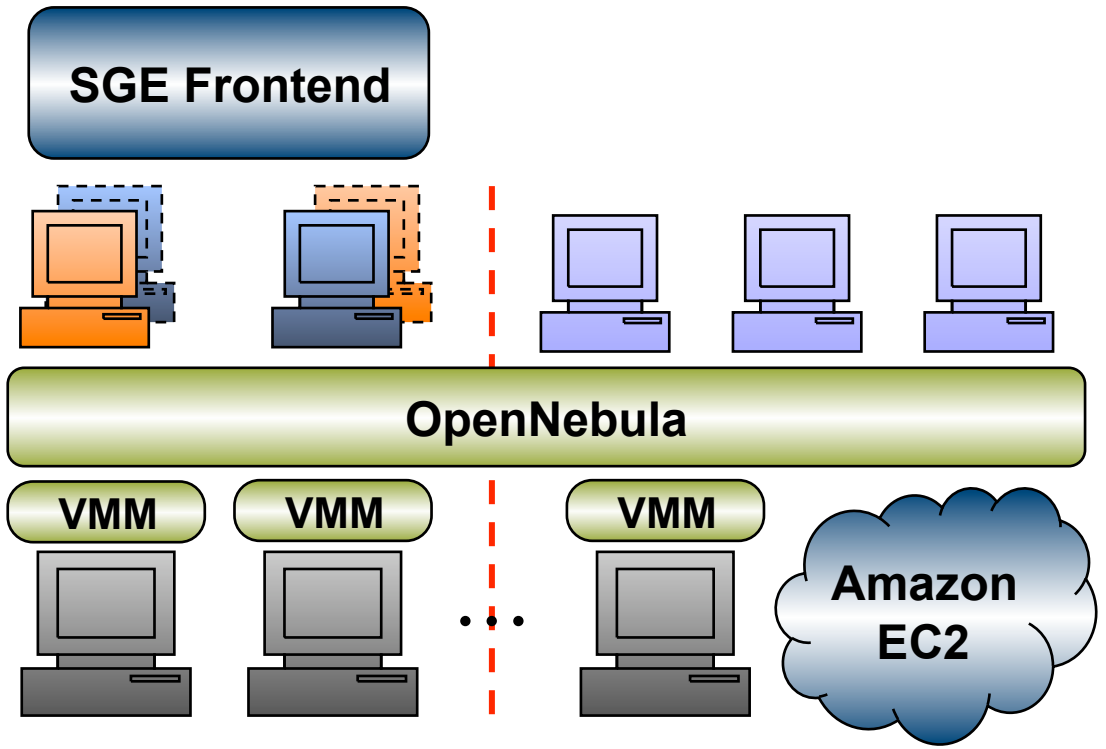

dsa-research.org

RESERVOIR Project

An Introduction to Virtualization and Cloud Technologies to Support Grid Computing

RESERVOIR

Who?

- IBM (coordinator), Sun, SAP, ED, TID, UCM, UNIME, UMEA, UCL, USI, CETIC, Thales and OGF-Europe
- 17-million and 3-year project partially funded by the European Commission (NESSI Strategic Project)

What?

- The Next Generation Infrastructure for Service Delivery, where resources and services can be **transparently and dynamically managed, provisioned and relocated like utilities** – virtually "without borders"

How?

- Integration of **virtualization technologies** with **grid computing** driven by new techniques for **business service management, driven by business use cases**

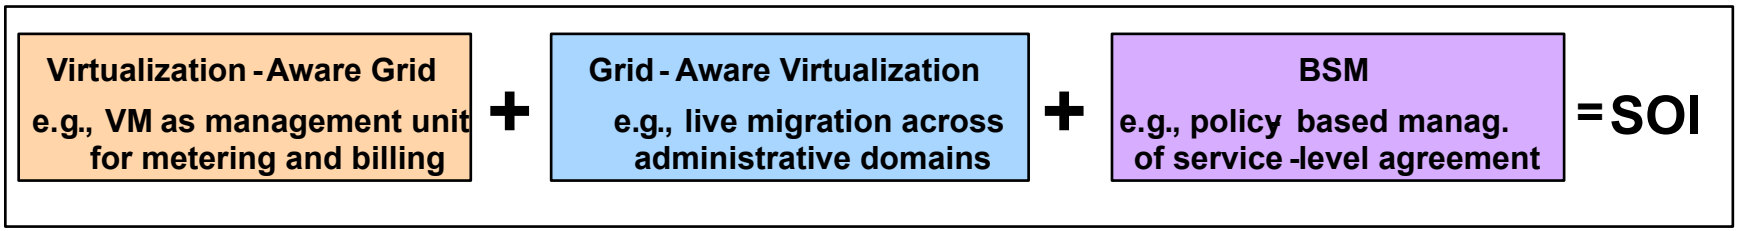

The Architecture, main Components and Interfaces

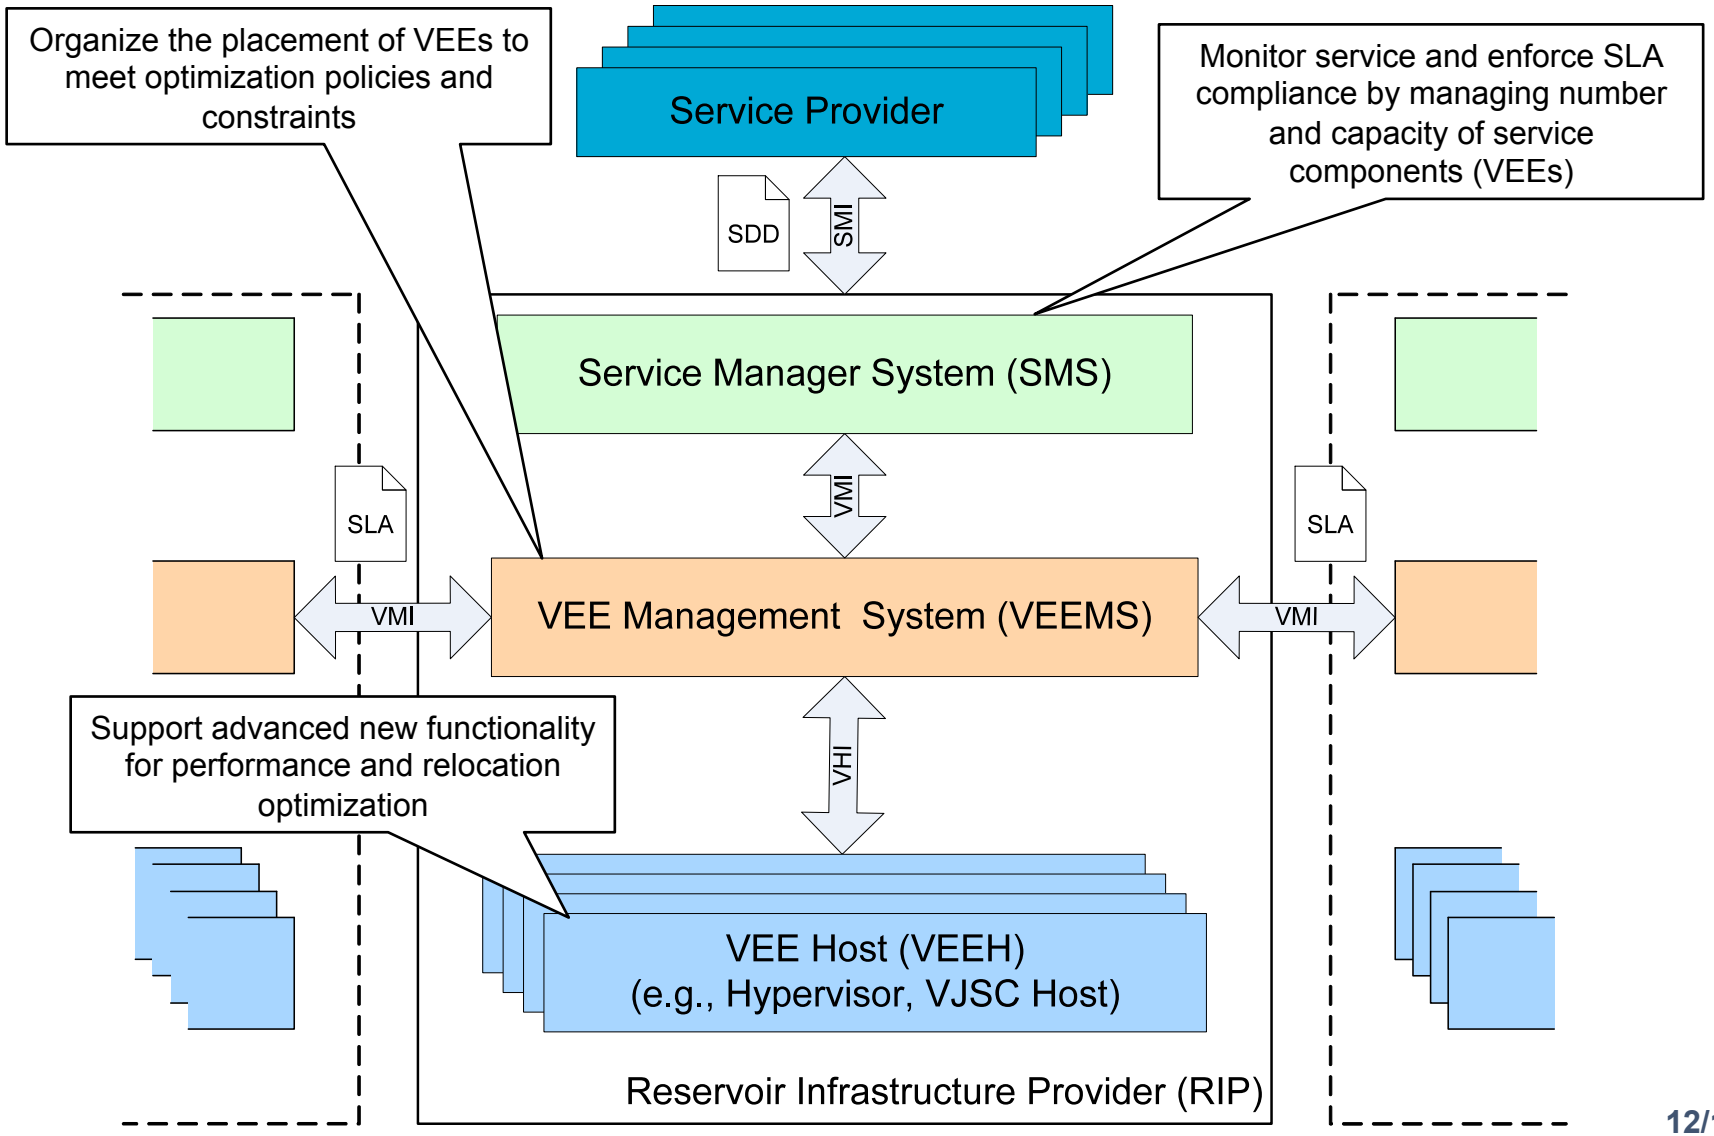

Conclusions

About the Coexistence of Grid, Virtualization and Clouds

- Virtualization, cloud, grid and cluster are **complementary technologies** and will coexist and cooperate at different levels of abstraction
- Virtualization and cloud **do NOT require any modification** within service layers from both the administrator and the end-user perspectives
- **Separation between service and infrastructure layers** will allow the application of the utility model to Grid/cluster/ HPC computing

THANK YOU FOR YOUR ATTENTION!!!
More info, downloads, mailing lists at
www.OpenNebula.org

OpenNebula is partially funded by the "RESERVOIR– Resources and Services Virtualization without Barriers" project
EU grant agreement 215605

www.reservoir-fp7.eu/

Acknowledgements

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

Demo on Scaling-out Local Infrastructures