

# CISCO Cloud Computing Research Symposium (C<sup>3</sup>RS)

November 5 & 6, 2008 San Jose, CA

## Virtual Machine Management with OpenNebula in the RESERVOIR project

Ruben Santiago Montero

[dsa-research.org](http://dsa-research.org)

Distributed Systems Architecture Research Group  
Universidad Complutense de Madrid



- Describe the goals of the **RESERVOIR project** and provide an overview of its architecture and design principles
- Discuss the challenges of **managing VMs in a distributed environment** and present the VM management model adopted by OpenNebula
- Study the operation of a Distributed VM Manager in **clouds environments**
- A simple use case: **Elastic Management of Computing Clusters**

## 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?

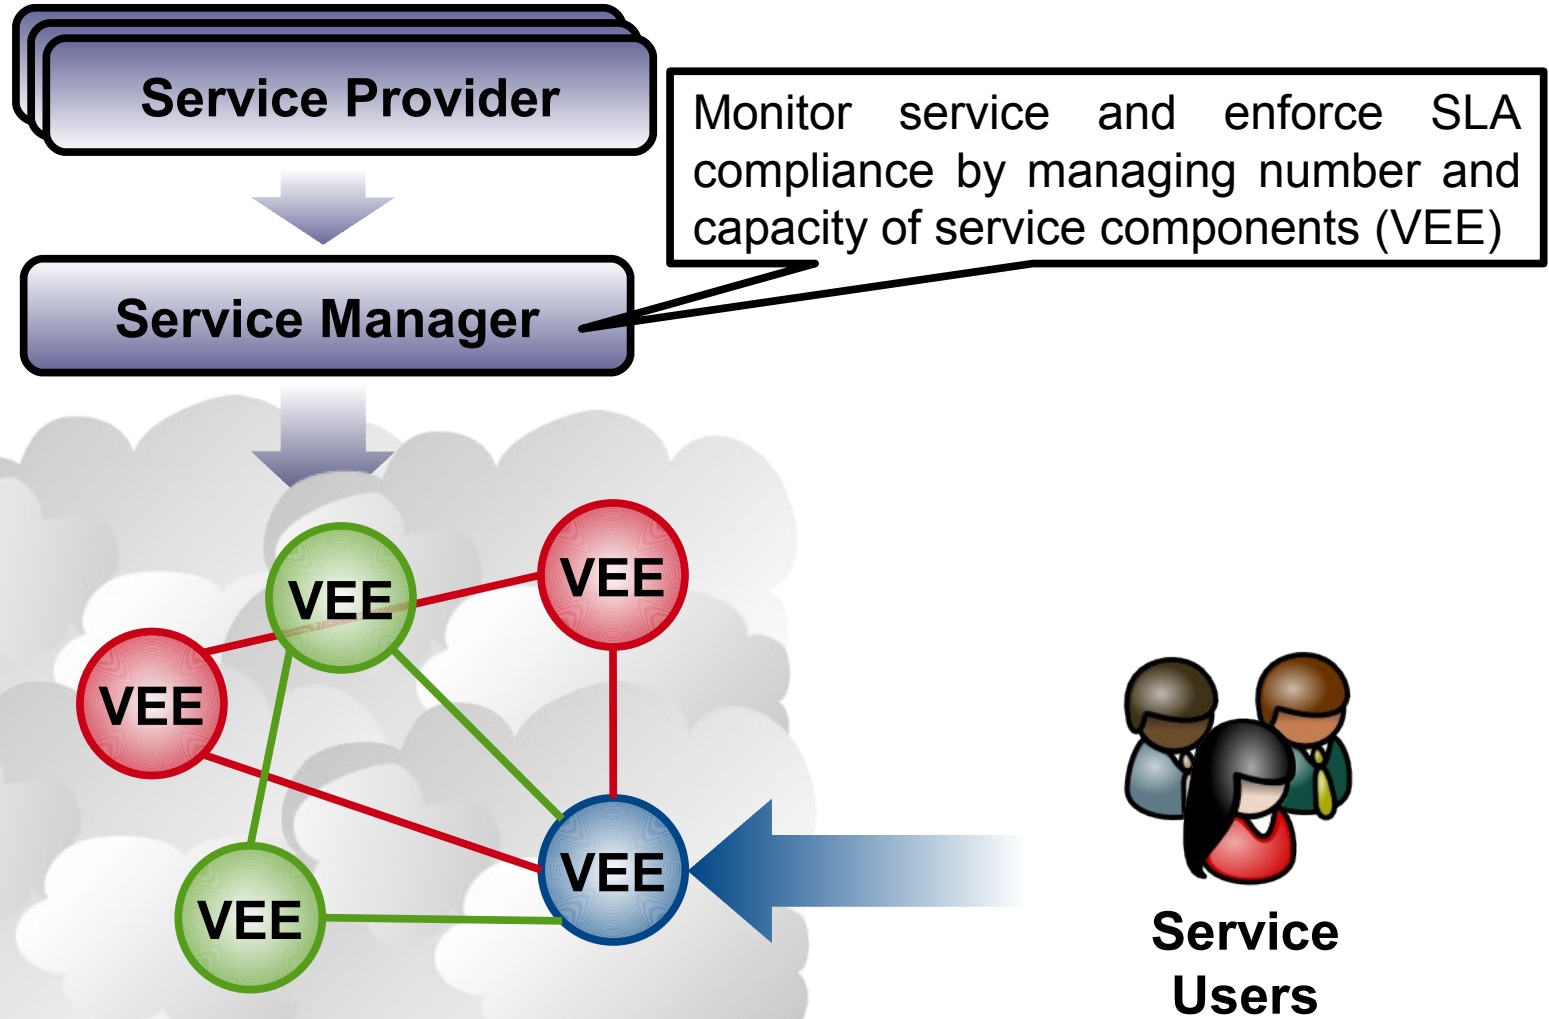
- The Next Generation Infrastructure for Service Delivery, where resources and services can be transparently and dynamically managed, provisioned and relocated like utilities – virtually “without borders”

## How?

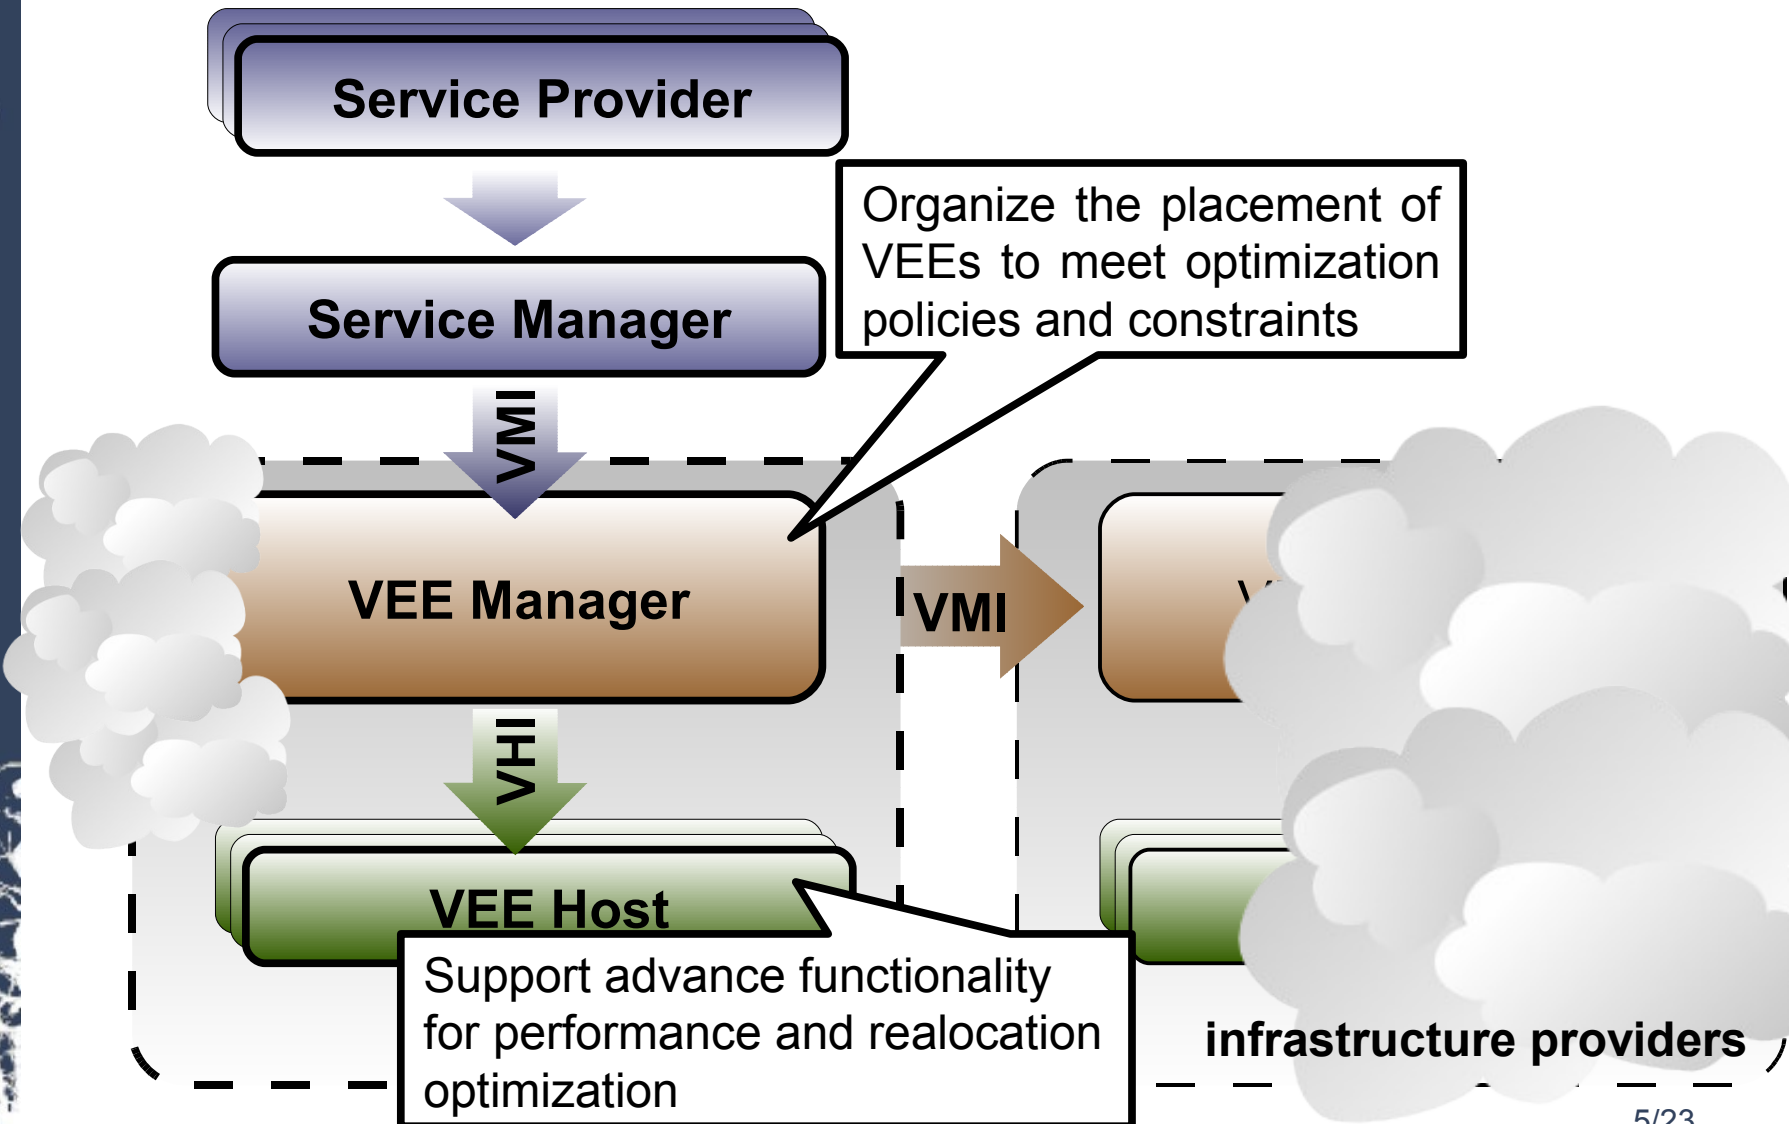
- Integration of virtualization technologies with grid computing driven by new techniques for business service management



## The Architecture, main Components and Interfaces

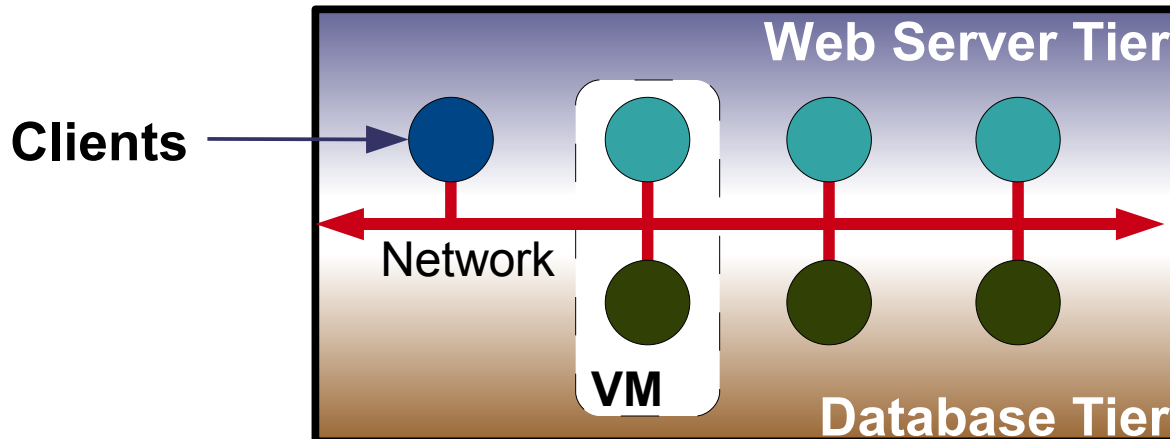


## The Architecture, main Components and Interfaces




## Virtual Machine Model

- The **service** as a first-class management entity
- Service **structure**
  - Service components run in VMs
  - Inter-connection relationship
  - Placement constraints
- The VM Manager is **service agnostic**
- However, it should provide **infrastructure context**



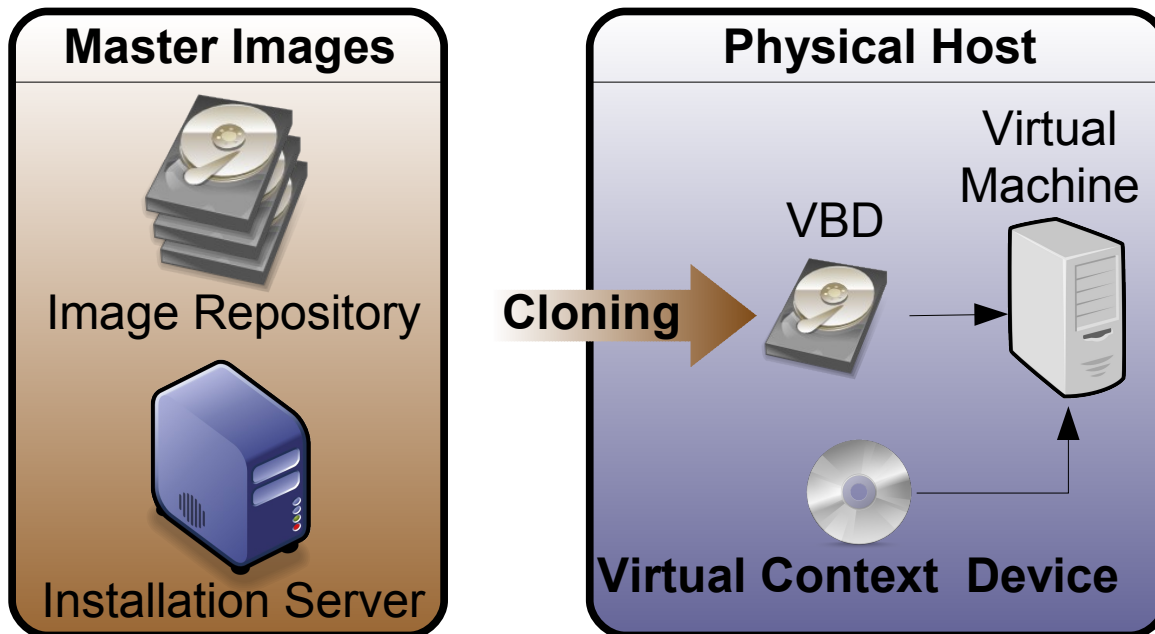
## Distributed Virtual Machine Management System

- Provides a uniform view of the resource pool
- Resources organized in a cluster architecture
- **Life-cycle management** and monitoring of VM
- The VM Management System **integrates** Image, Network and Virtualization technologies



## Image Management

- VM Images Sources:
  - Master images in local repositories
  - Appliance supplier
  - Creation on the fly
- Clones have to be contextualized (Context VBD)

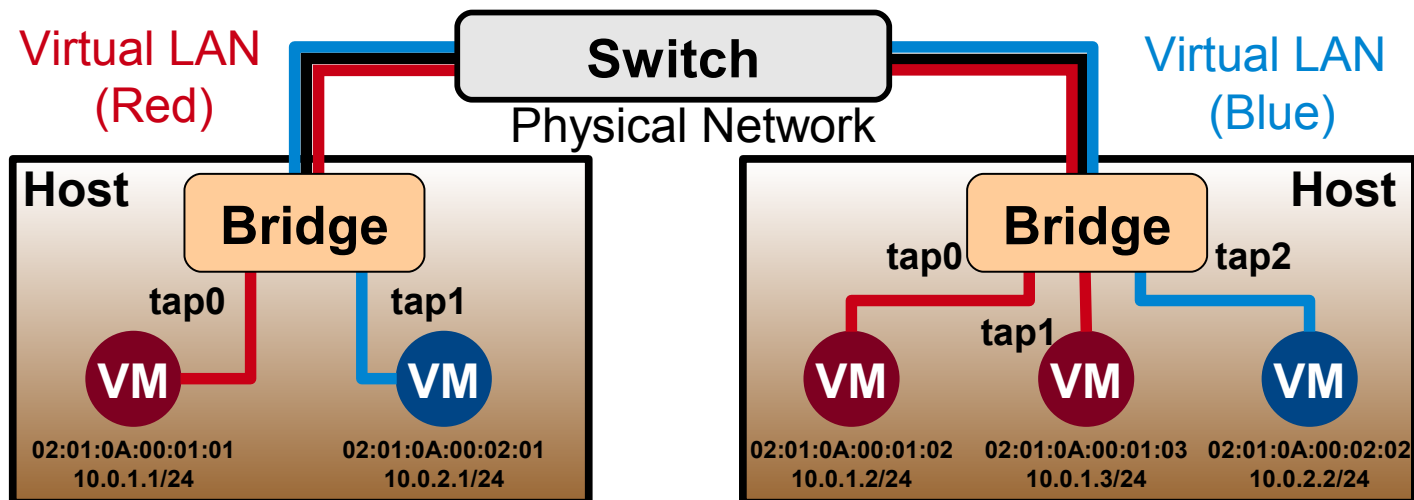




## Network Management

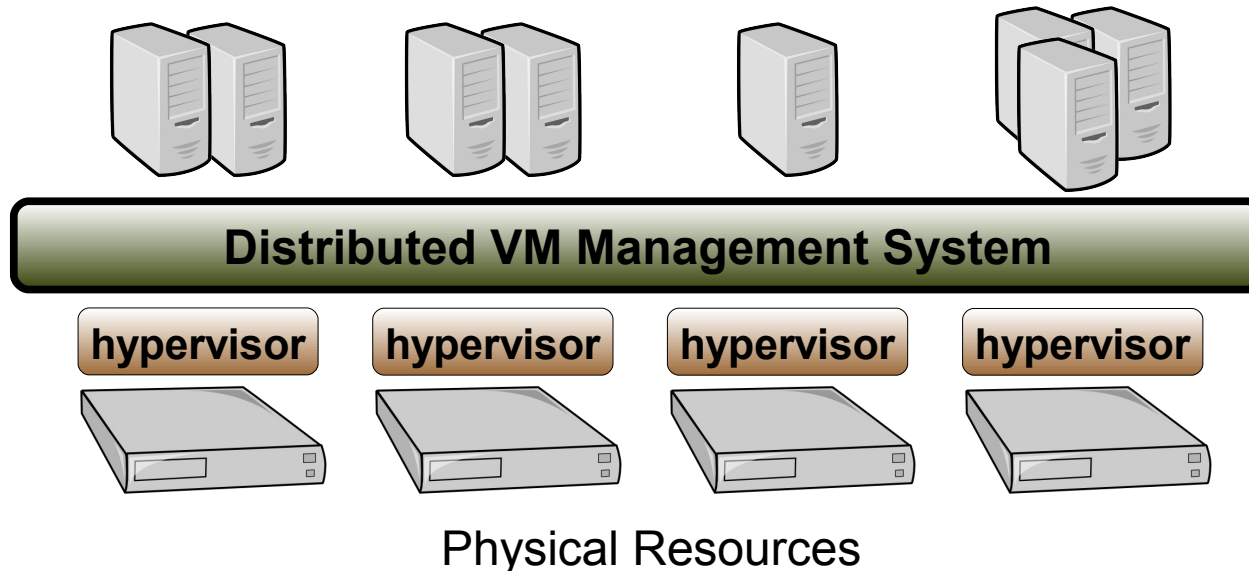
- VMs interconnected through **one or more** networks
  - Isolated, layer 2 LANs
  - Virtual networks are dynamically created
  - Medium size networks (x.x.x.x/20) with limited public IPs
- **TCP/IP services** are not responsibility of the VM Manager

## Sample Implementation



## Virtualization

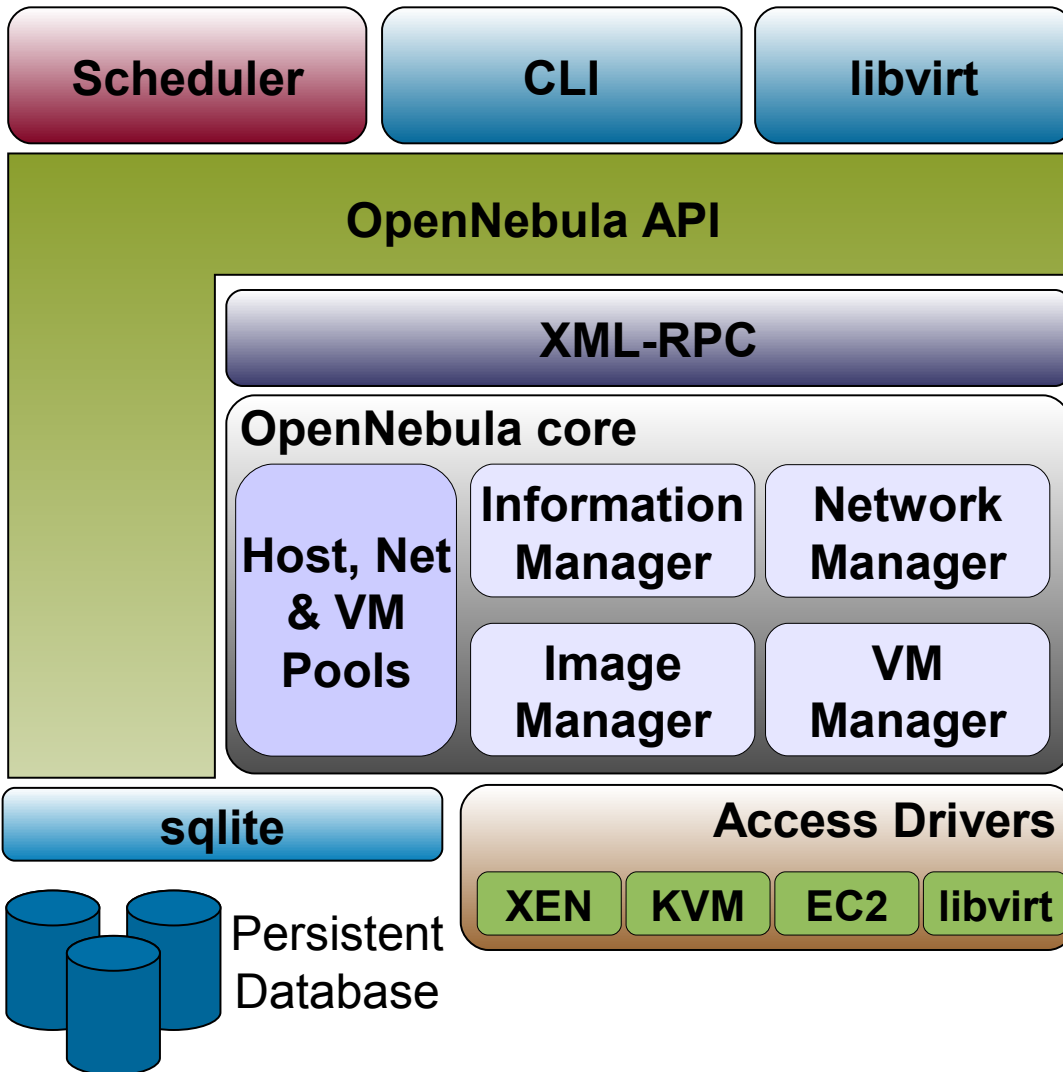
- Virtual Machine structure
  - One or more NICs attached to virtual or public networks
  - One or more system images (clones)
  - A context virtual block device
  - A **required capacity (memory, CPU)**
- Use the host hypervisor to create, monitor and control VMs



## Life-cycle of a Virtual Machine (or a set of)

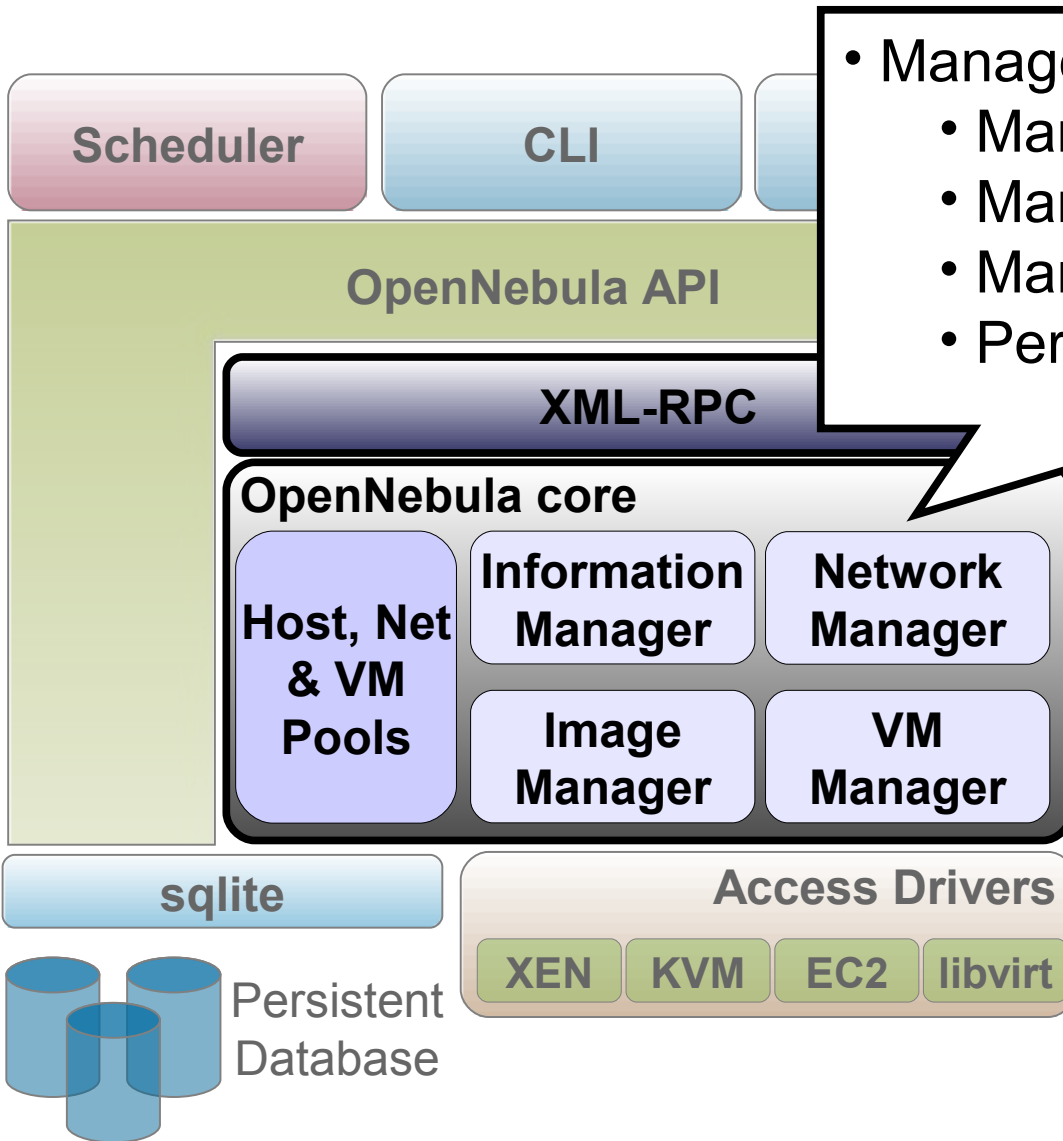
- **Resource Selection:** *Where do I place the VM?*
  - Capacity planning (consolidation)
  - Placement requirements (e.g. affinity)
  - Placement Heuristics (e.g. Green IT, AR...)
- **Resource Preparation:** *What do I need for the VM?*
  - Network preparation
  - Image cloning & contextualization
- **VM Creation:** *How do I start a VM?*
  - Interface with different hypervisors
- **VM Monitoring:** *How is the VM doing?*
- **VM Migration:** *Is there a better resource for the VM?*
  - Adjust placement to better fit to the infrastructure target
- **VM Termination:** *Do I need to save any VM image?*

## OpenNebula: Distributed VM Management System



- **Open Source** - Apache 2
- **Flexible & Open Design**
  - Third-party components
  - Easily adapted
  - Easily extended
- Different Hypervisors
- Operation in **Federated Environments**
- **Integral management** of Virtual Services

## OpenNebula: Distributed VM Management System

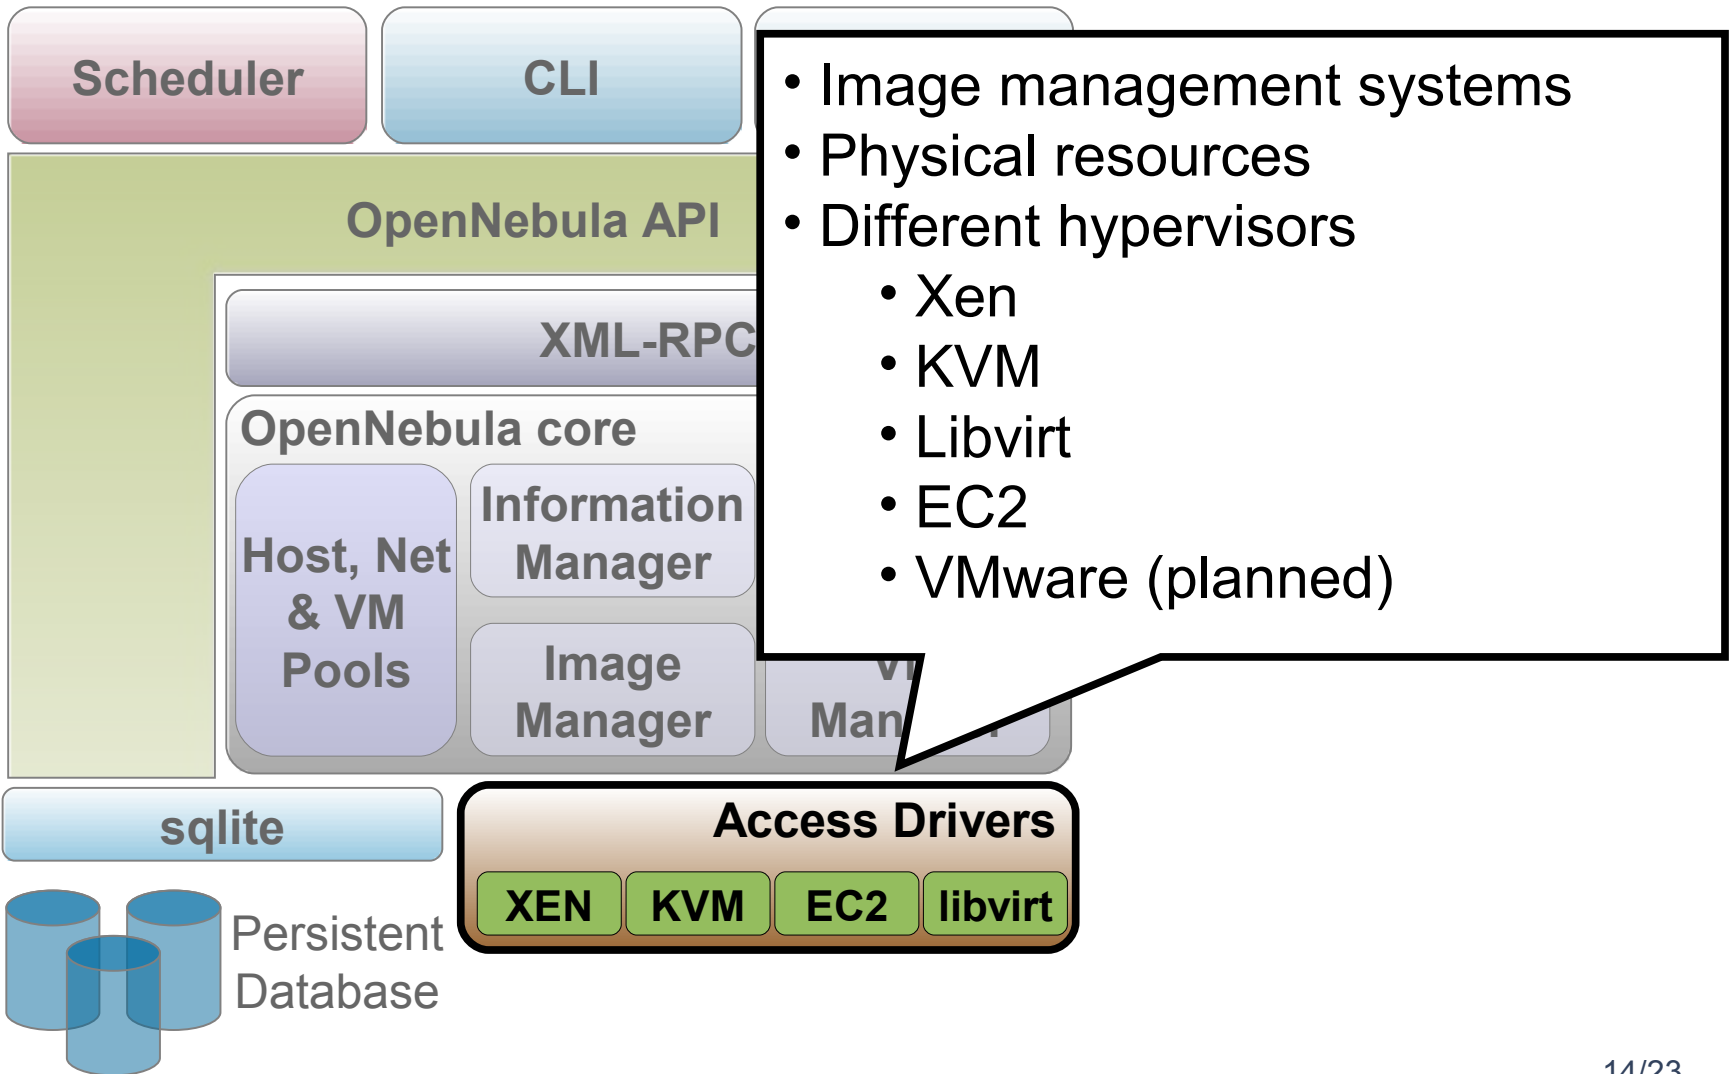


- Manage IP and MAC addresses
  - Manage VM life-cycle
  - Manage Image transfers
  - Manage Physical Resources
  - Persistent back-end

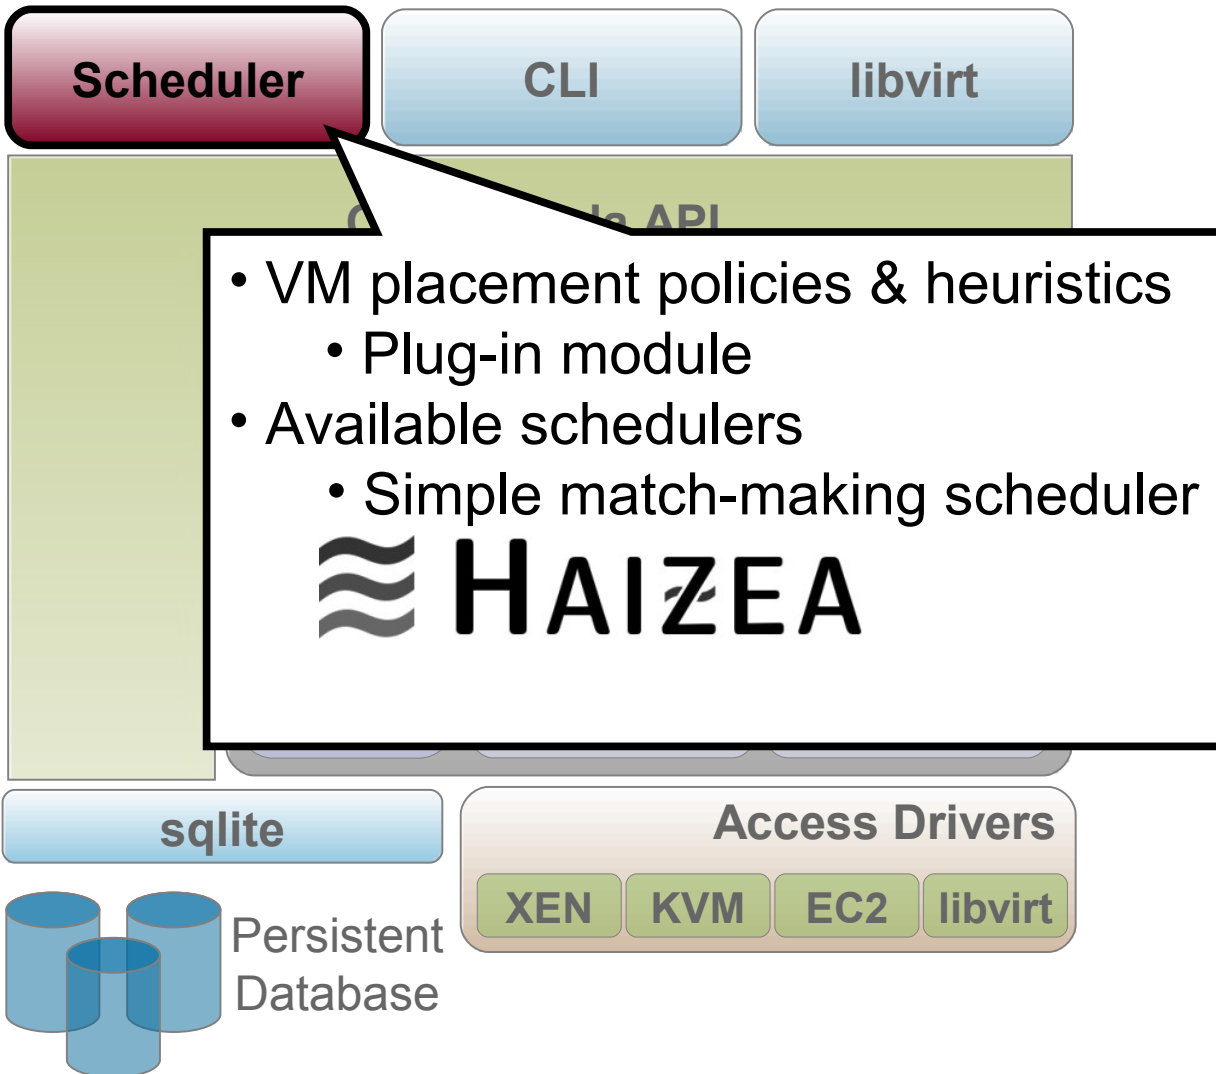
# Distributed Management of Virtual Machines

Virtual Machine Management with OpenNebula in the RESERVOIR project

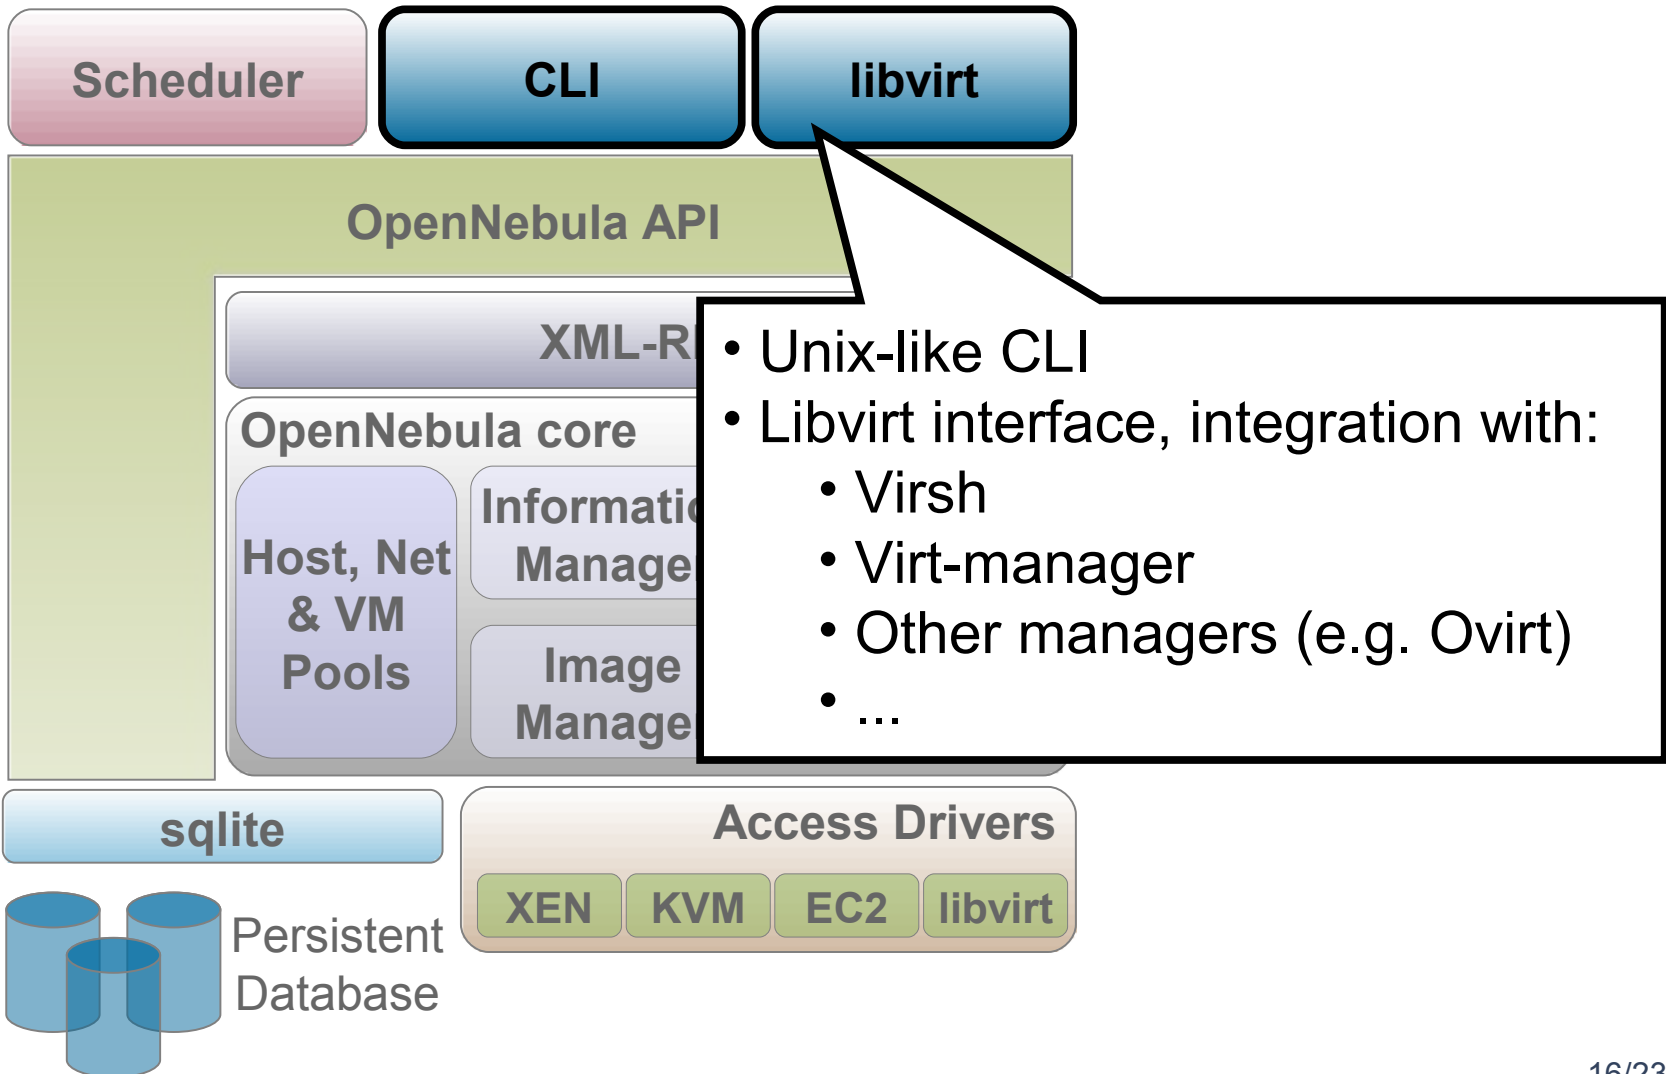
## OpenNebula: Distributed VM Management System



## OpenNebula: Distributed VM Management System



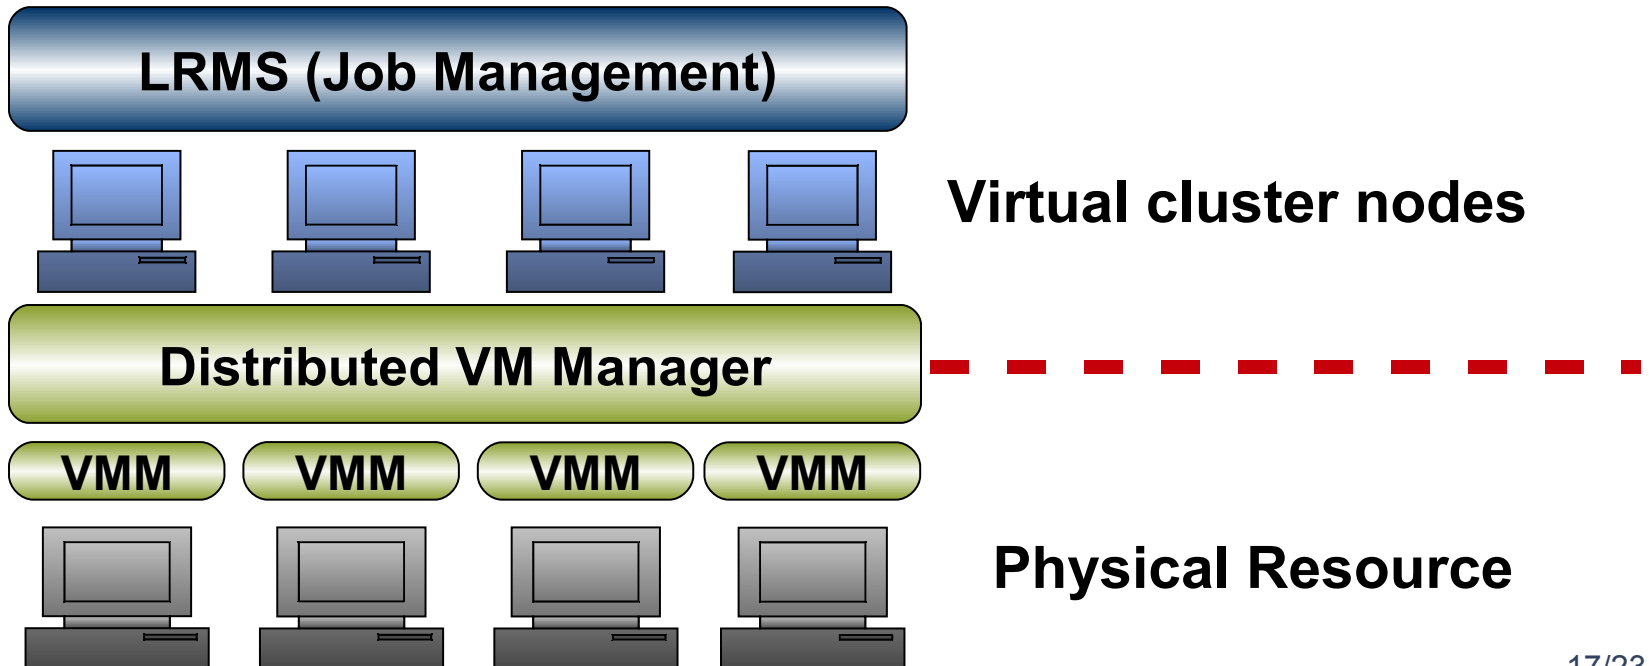
## OpenNebula: Distributed VM Management System





## A New Infrastructure Layer...

- Separation of Resource Provisioning from Job Management
- Seamless integration with the existing middleware stacks.
- Completely transparent to the computing service and end users




# Elastic Management of Computing Clusters


*Virtual Machine Management with OpenNebula in the RESERVOIR project*

**Cluster users**

**HTTP clients**




**Virtual Network**



**Cluster  
Front-end**






**Virtual Workernodes**



**Web  
Server**

**Service Layer**

**OpenNebula**



**Physical Infrastructure**

**Infrastructure Layer**

# Elastic Management of Computing Clusters

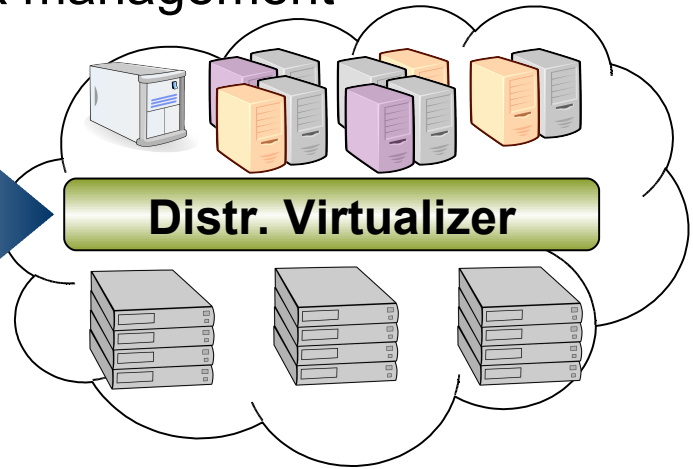

Virtual Machine Management with OpenNebula in the RESERVOIR project

## A Service to Provide Hardware on Demand (IaaS)

- Cloud systems provide **virtualized resources as a service**
- Provide remote on-demand access to infrastructure (through VMs)
- Main components of a **Cloud architecture**:
  - **Front-end**: Remote interface
  - **Back-end**: Local VM, image & network management

**Simple Interfaces**

- VM Management
- Image Management



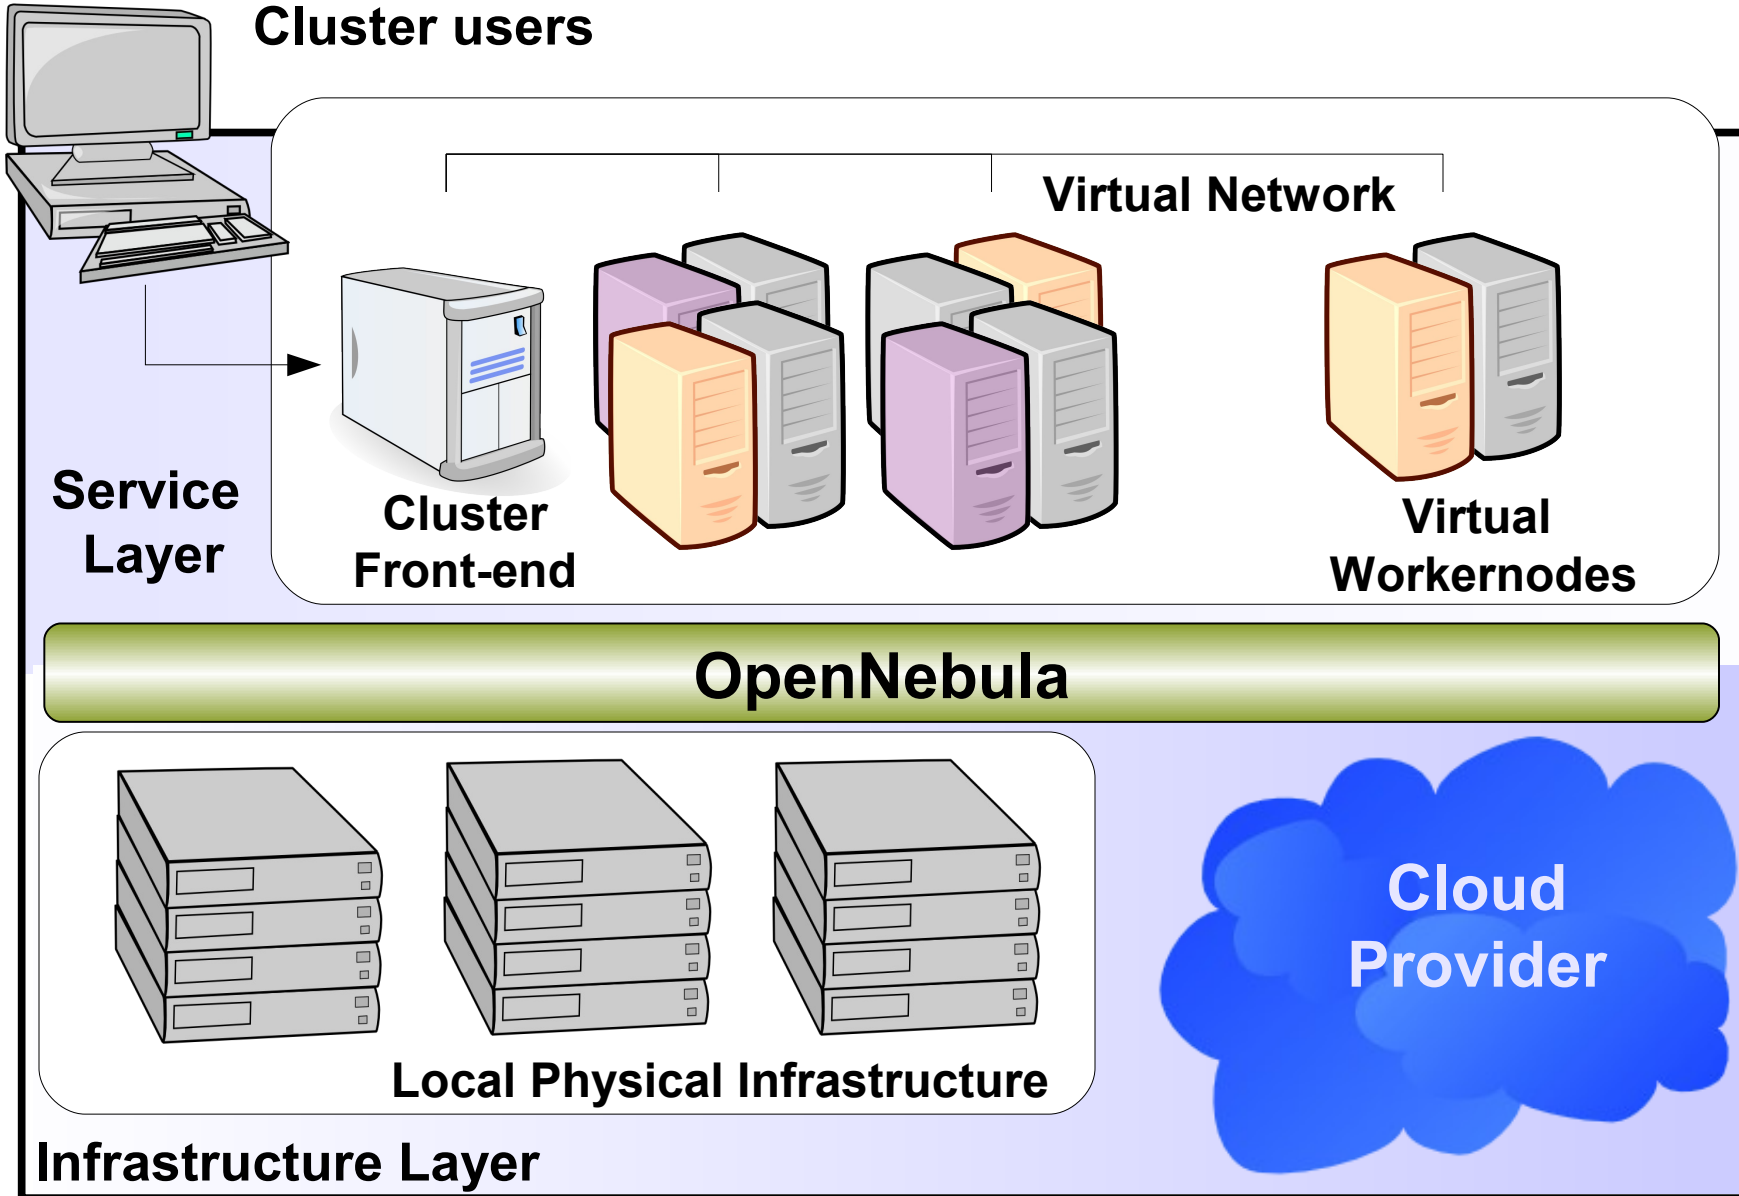
## Infrastructure Cloud Services

- Commercial Cloud: Amazon EC2, GoGrid...
- Scientific Cloud: Nimbus (University of Chicago)

# Elastic Management of Computing Clusters

*Virtual Machine Management with OpenNebula in the RESERVOIR project*

dsa-research.org

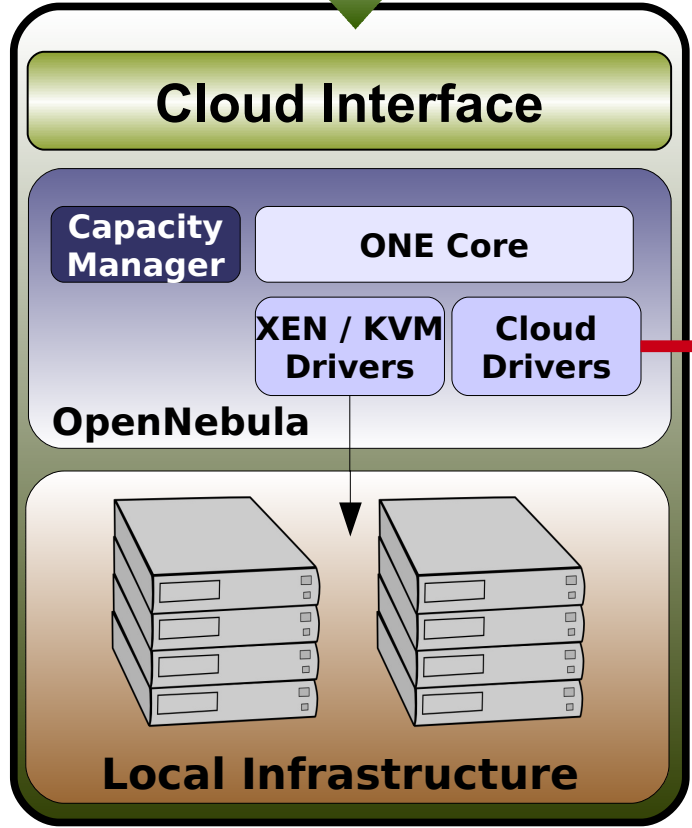


# Federation of Cloud Sites

Virtual Machine Management with OpenNebula in the RESERVOIR project

Clients or other Provider

- Consumer/Provider Relationships
- Nimbus integration (WSRF + EC2)
- RESERVOIR Cloud VMI
- Eucalyptus (planned)



- Federation of Cloud Providers
- Image, Network and VM issues





THANK YOU FOR YOUR ATTENTION!!!

More info, downloads, mailing lists at  
[www.OpenNebula.org](http://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/](http://www.reservoir-fp7.eu/)

## Acknowledgements

---

- Javier Fontan
- Tino Vazquez
- Ignacio M. Llorente
- Rafael Moreno



# Elastic Management of Computing Clusters (DEMO - BACKUP)

dsa-research.org

