Danish Grid Forum – Minisymposium on Cloud Computing
May 12th, Niels Bohr Institute, Copenhagen

Clouds and Virtualization to Support Grid Computing

Ruben Santiago Montero

dsa-research.org

Distributed Systems Architecture Research Group
Universidad Complutense de Madrid

Objectives

- Brief review of some limitations of current Grids
- Use of virtual machines in Grids and its use for the dynamic provisioning of virtual clusters
- Grids & Clouds: Scale-out a Grid site
- Demo!

Some Limitations of Current Grids

- High degree of heterogeneity (software & hardware)
- High operational costs
- Isolate and partition resources contributed to the Grid
- Specific environment requirements for different Vos
- Users simply do not feel like adopting our execution models (*pilot jobs*...)

Grids are difficult to maintain, operate and use

Virtual Machines

- A VM is an isolated runtime environment (guest OS and apps)
- Hypervisors: Full Virtualized, para-virtualization, HW Virtualization

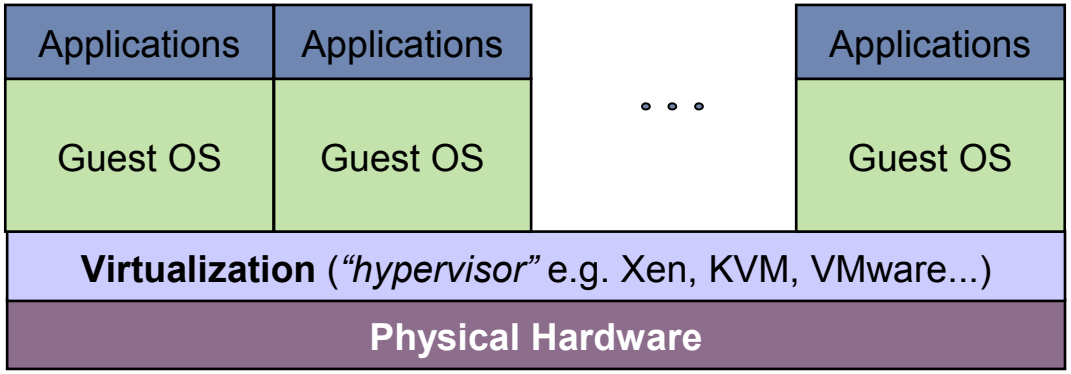

Benefits of Virtualization Platforms

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

Virtual Infrastructure Manager (VIM)

- ...but something more is needed
  - Where did/do I put my VM? (***scheduling & monitoring***)
  - How do I provision a new cluster node? (***clone & context***)
  - What MAC addresses are available? (***networking***)
- Provides a ***uniform view*** of the resource pool
- ***Life-cycle management*** and monitoring of VM
- The VIM ***integrates*** Image, Network and Virtualization

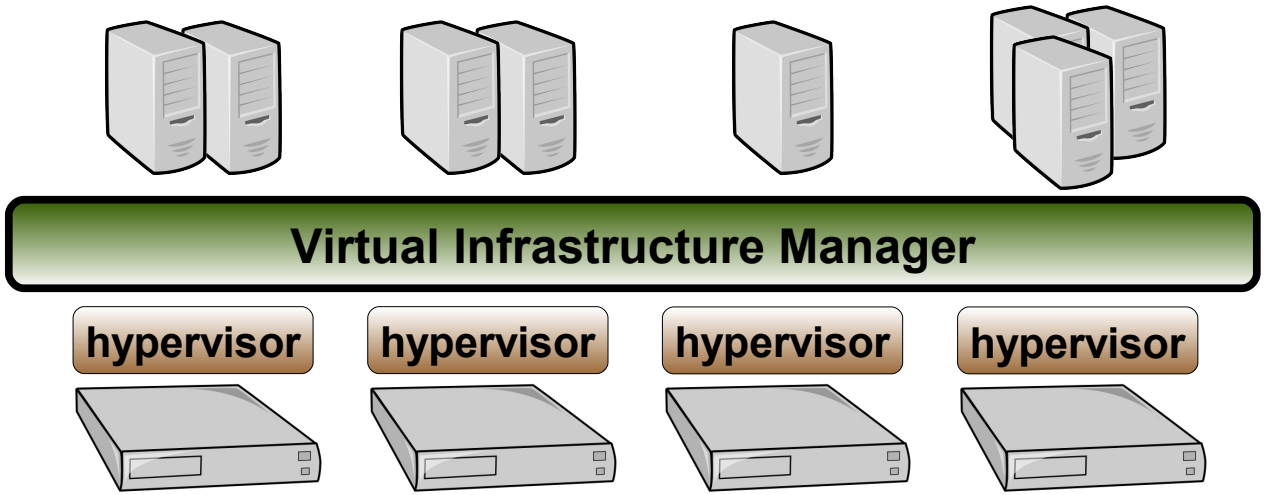

The OpenNebula Virtual Infrastructure Manager

www.OpenNebula.org

- Flexible & Open Design
  - Third-party components
  - Easily adapted & extended
- Management of *Virtual Services*
  - Image, Network & Context
- Integrated with cloud providers
- Open Source – Apache2
- Included in Ubuntu 9.04 (server)

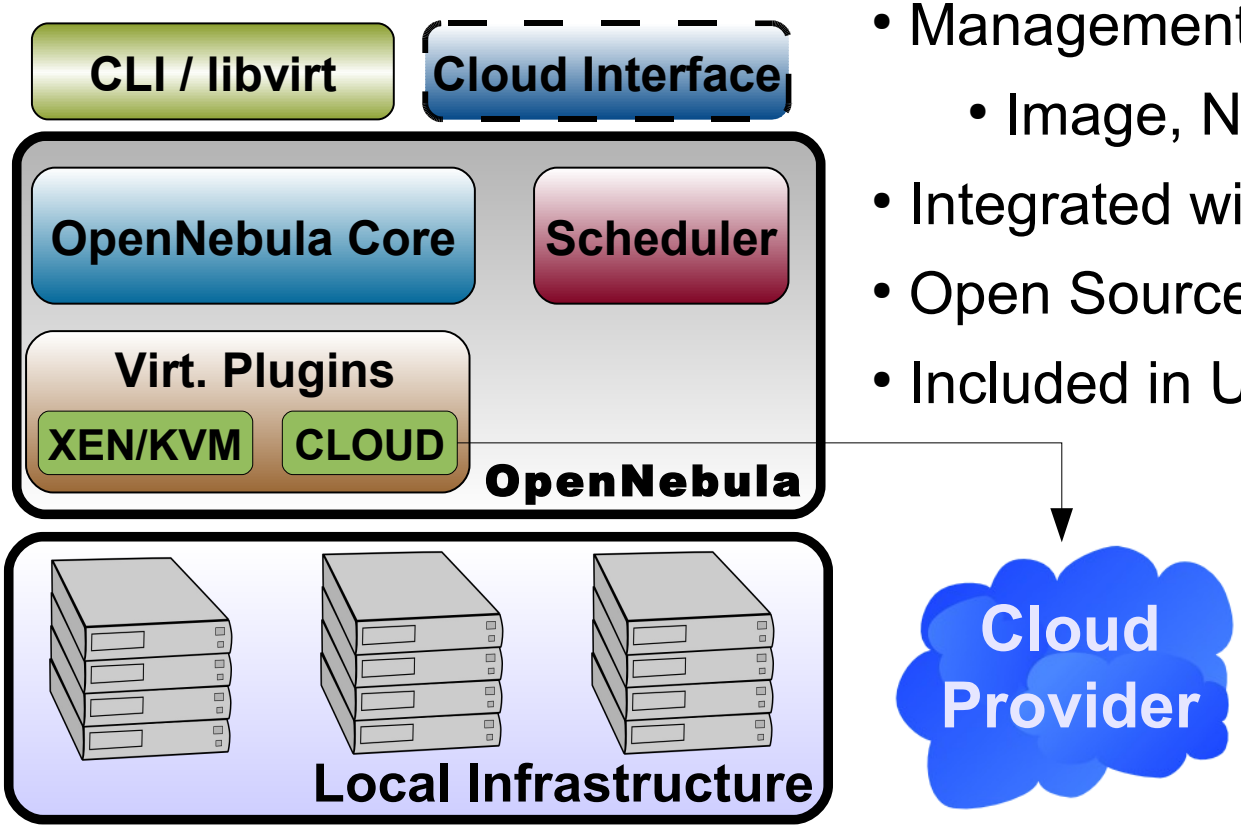

Grids & Virtual Machines

New provision models for Grids: Virtualization and Clouds

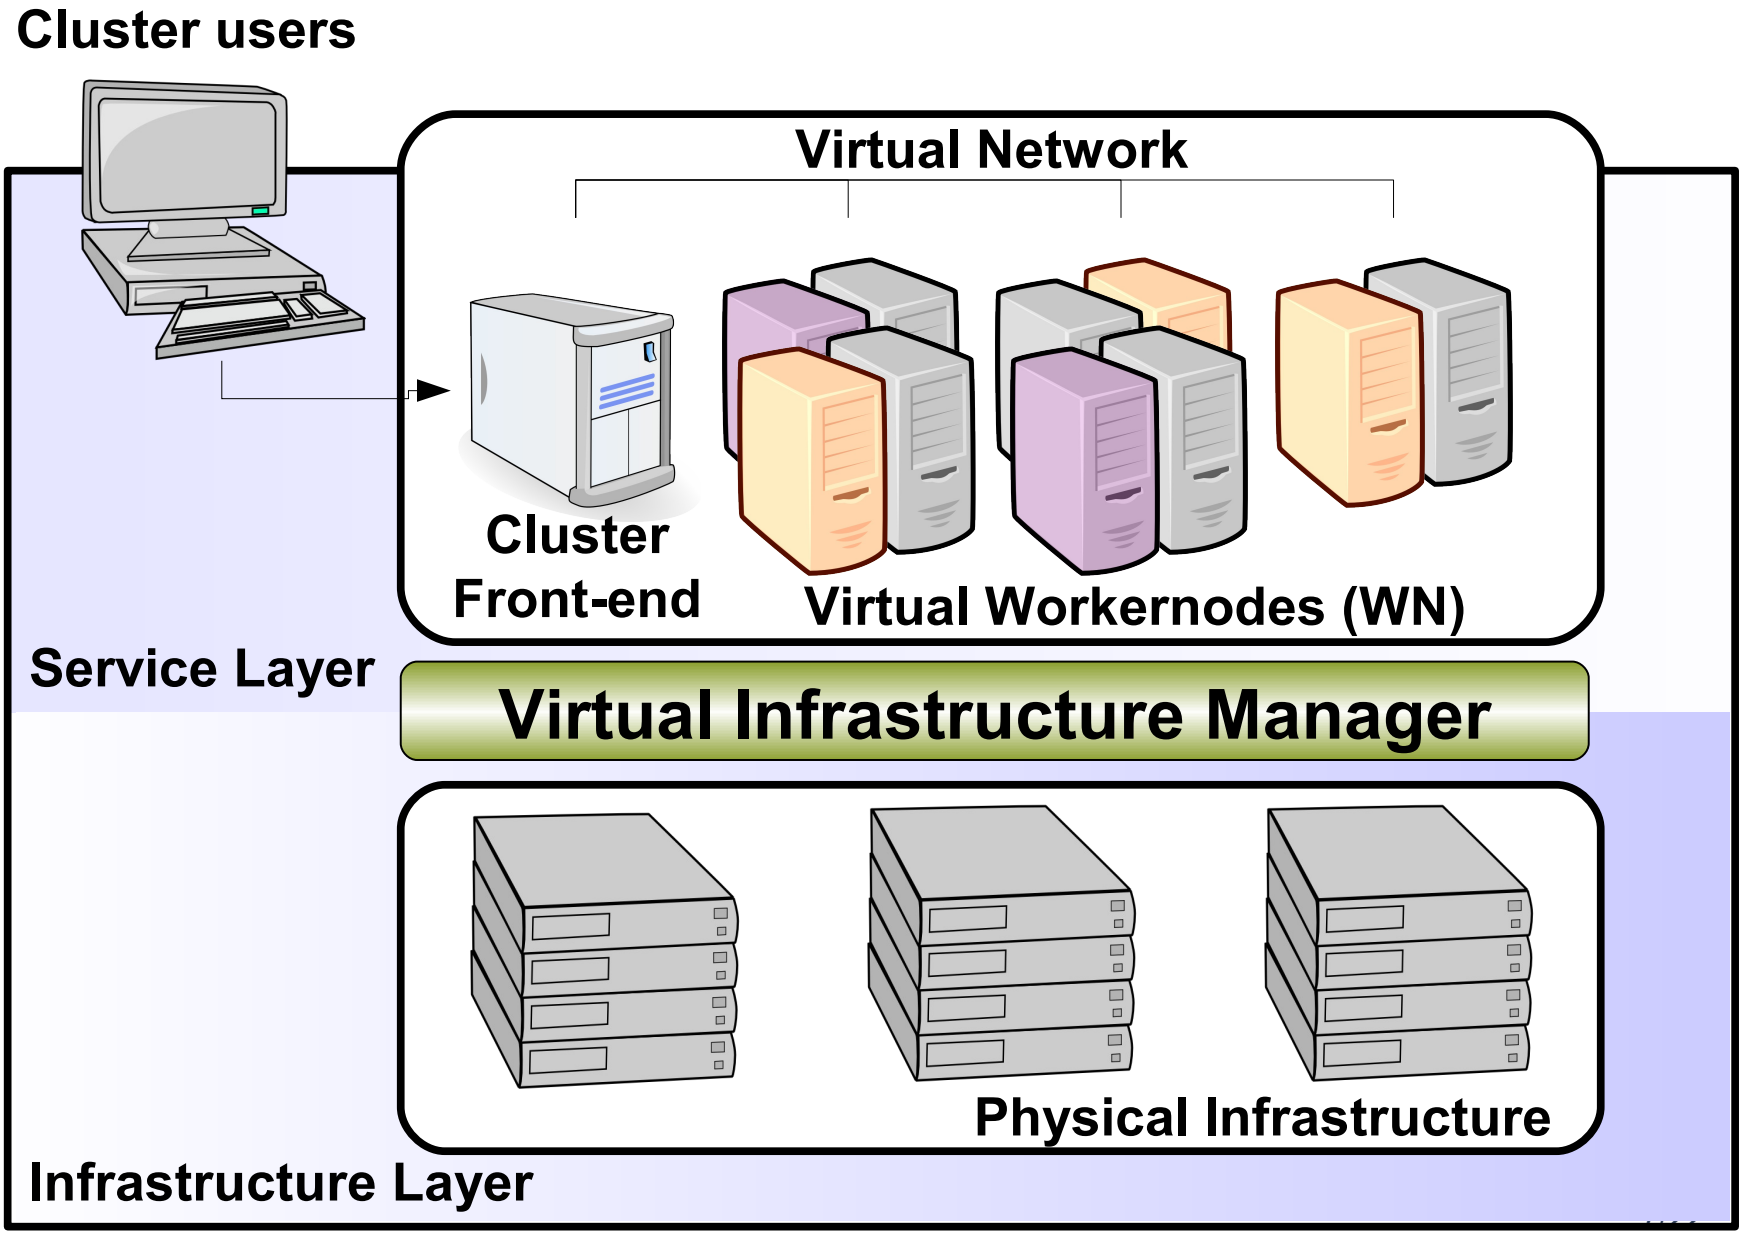

Grids & Virtual Machines

New provision models for Grids: Virtualization and Clouds

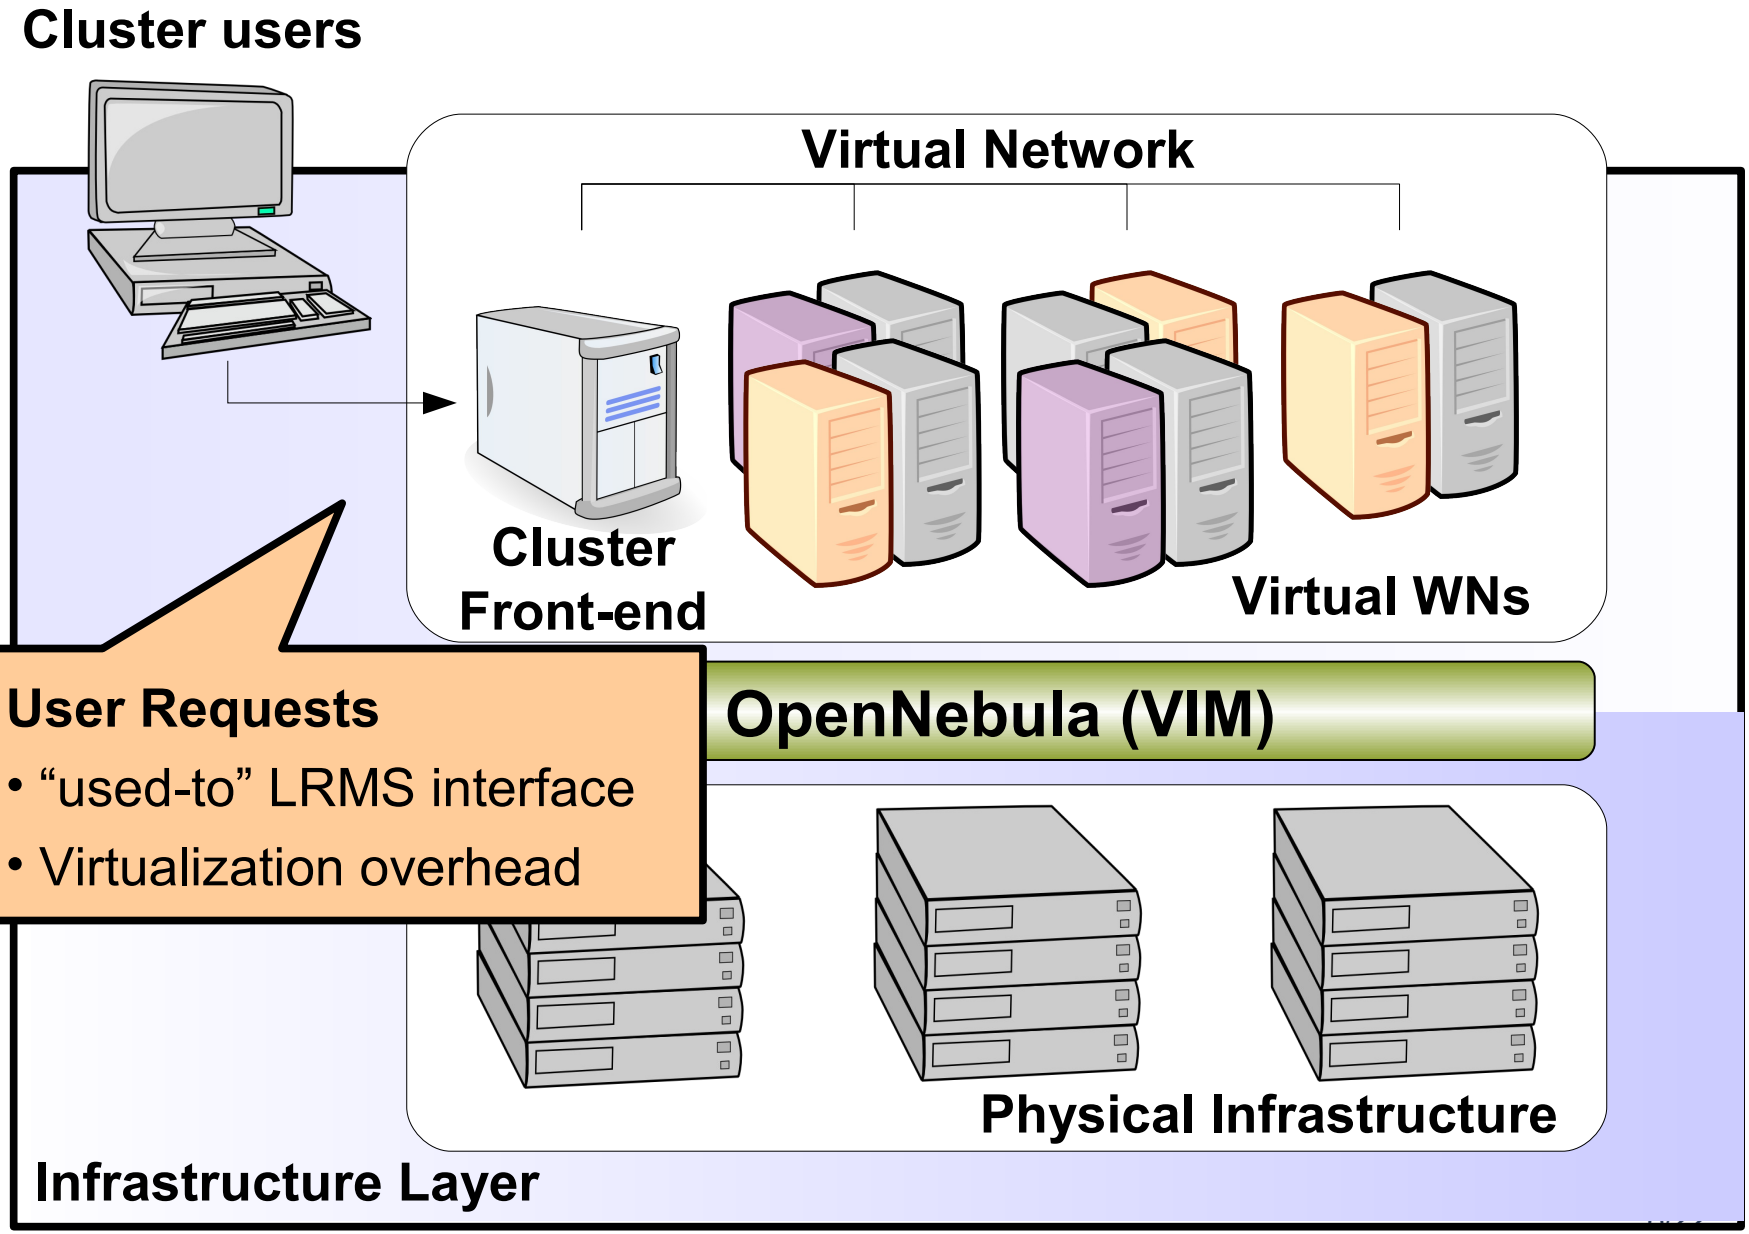

Grids & Virtual Machines

New provision models for Grids: Virtualization and Clouds

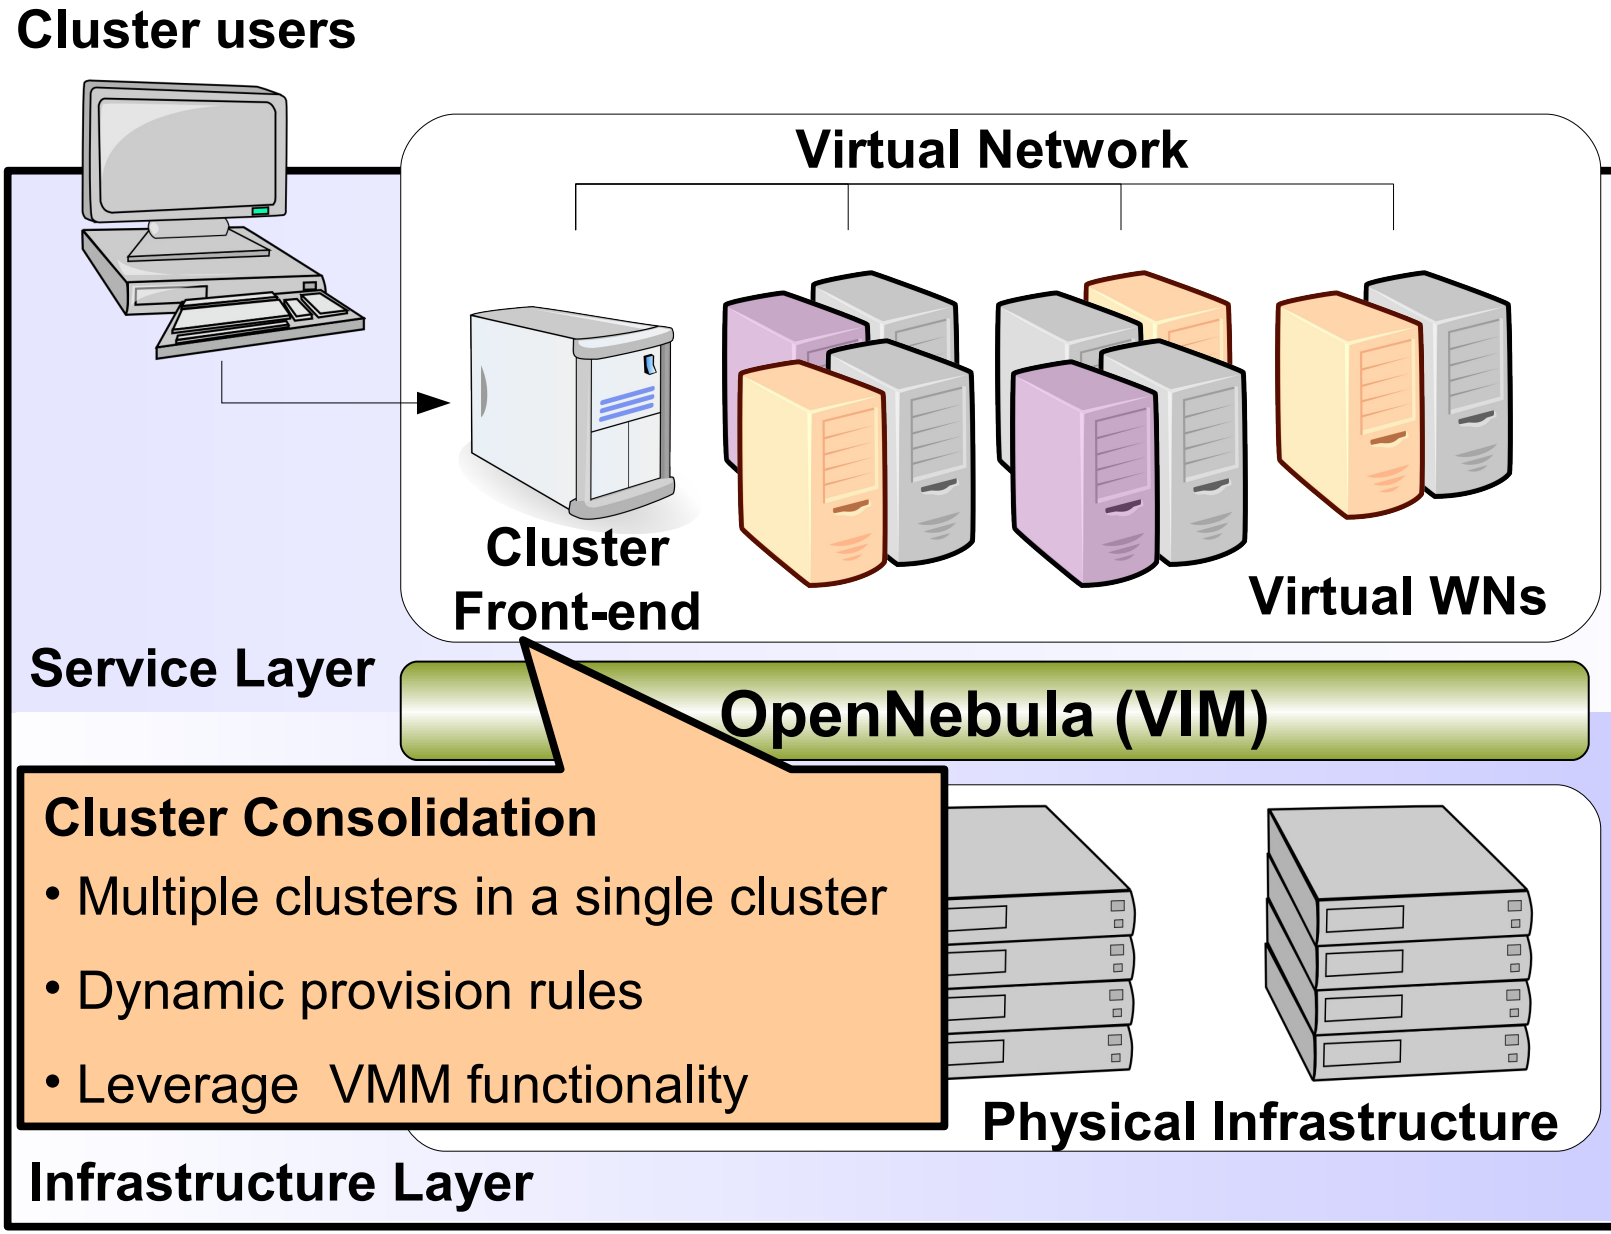

Grids & Virtual Machines

New provision models for Grids: Virtualization and Clouds

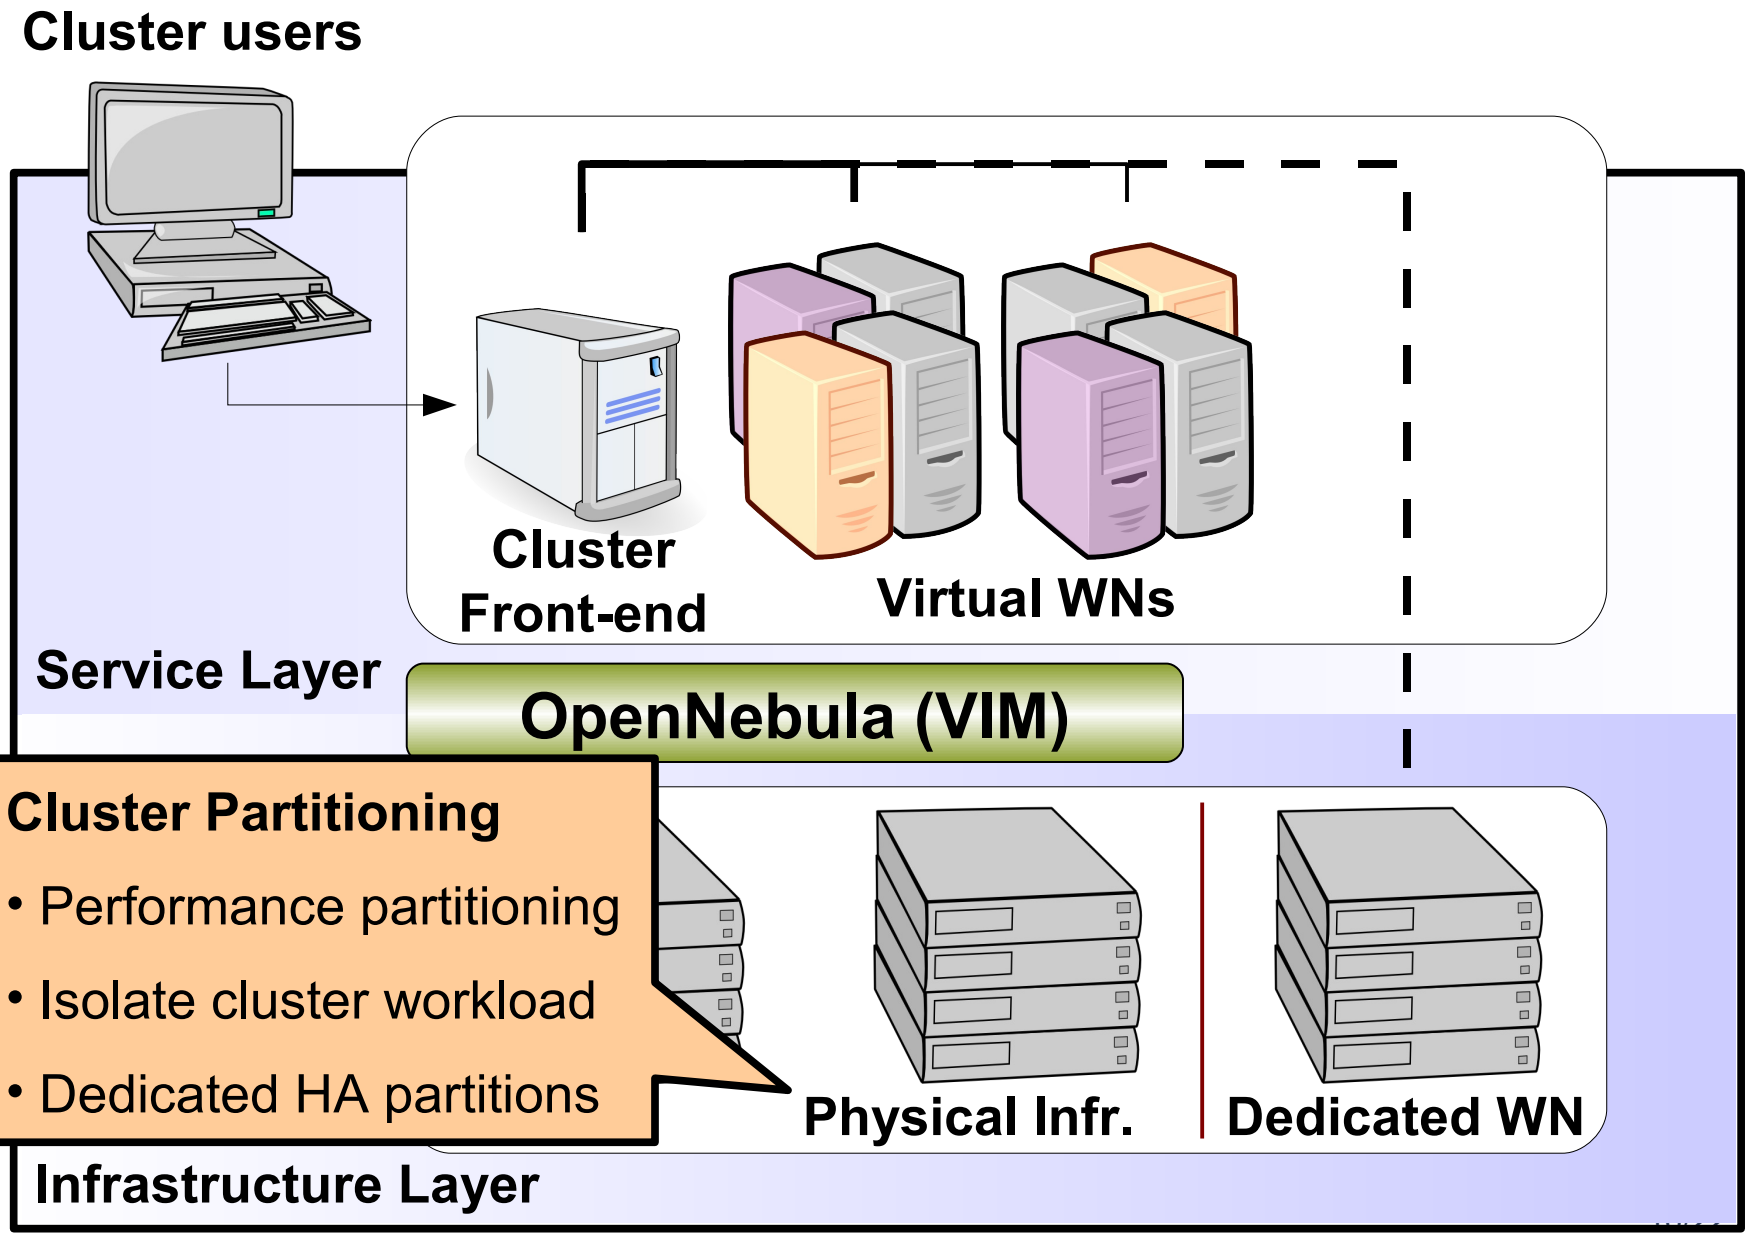

Grids & Virtual Machines

New provision models for Grids: Virtualization and Clouds

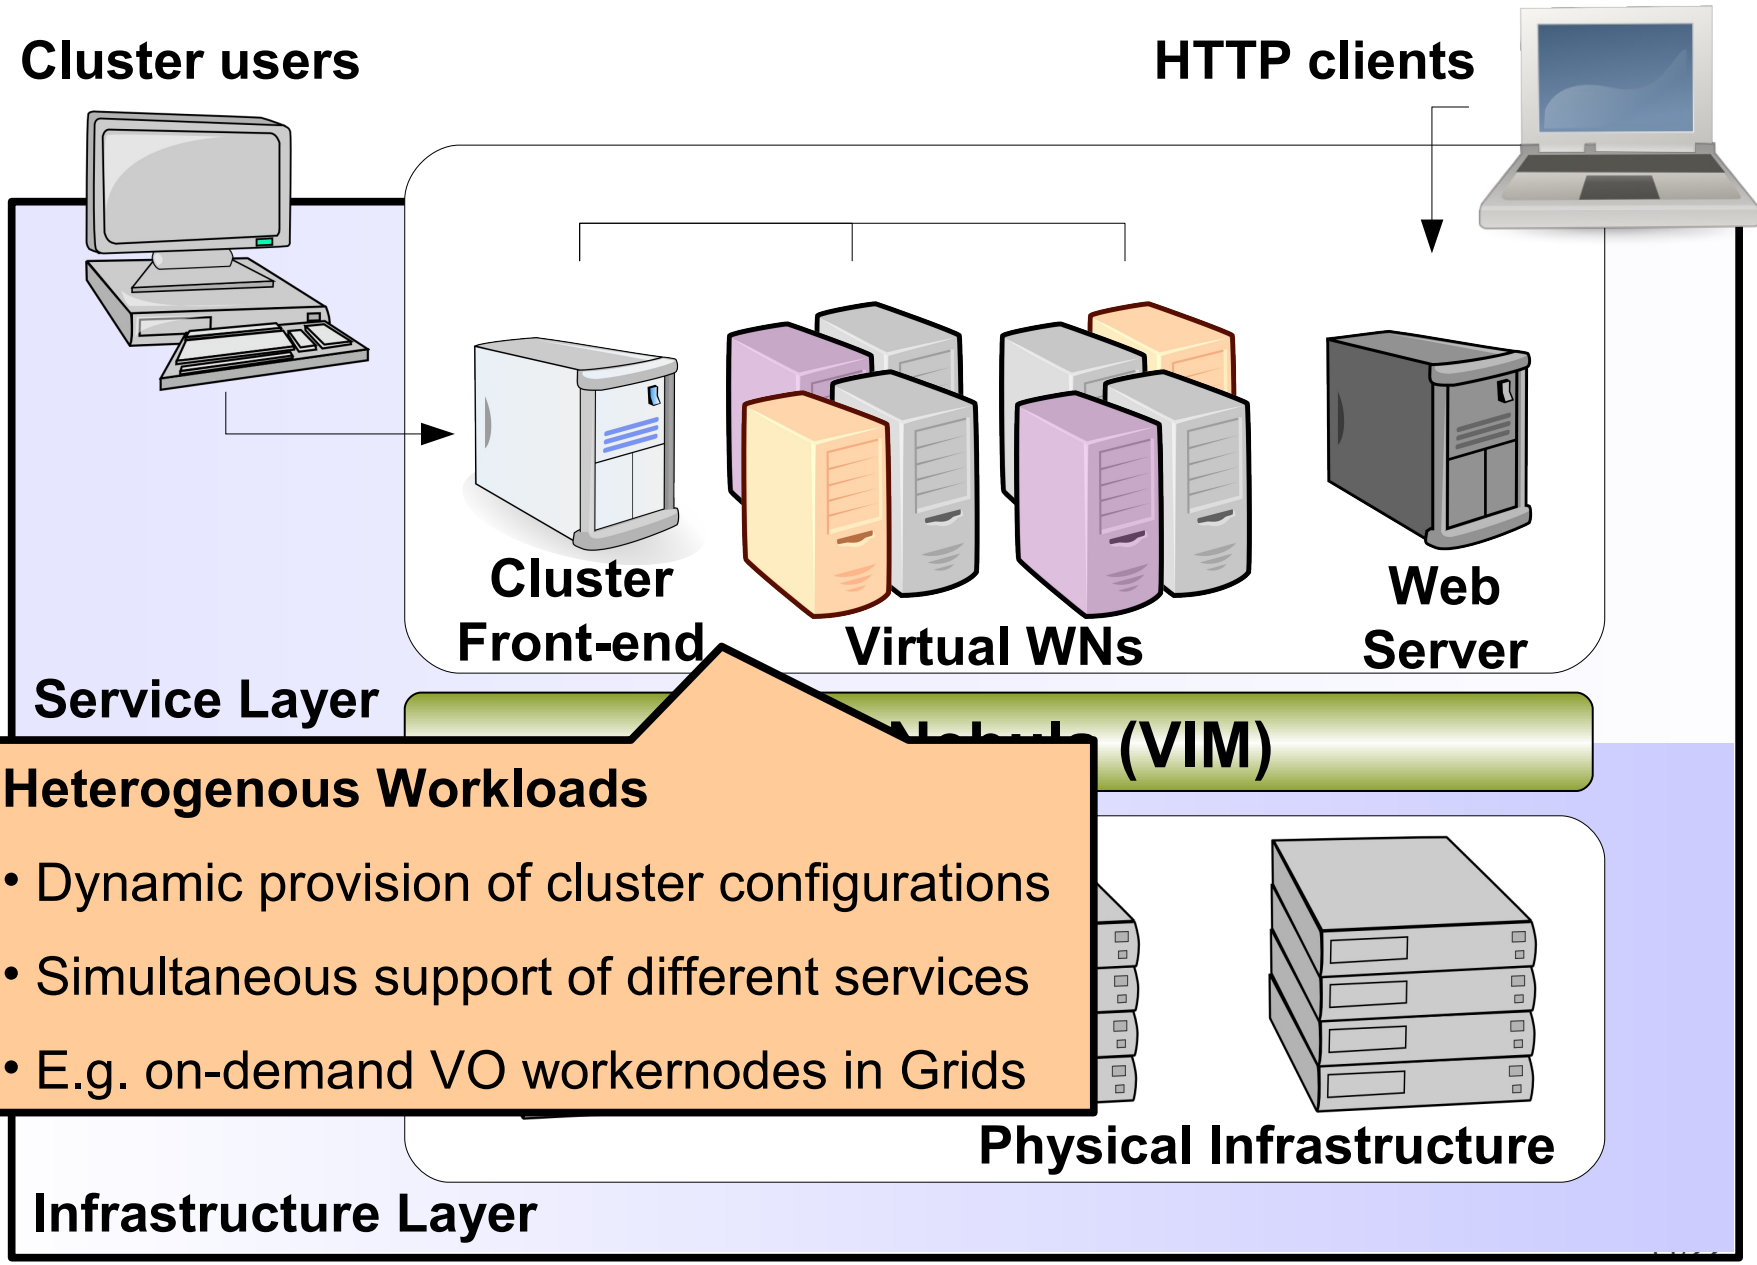

Heterogenous Workloads

- Dynamic provision of cluster configurations
- Simultaneous support of different services
- E.g. on-demand VO workernodes in Grids

A Complete Grid Middleware Stack

New provision models for Grids: Virtualization and Clouds

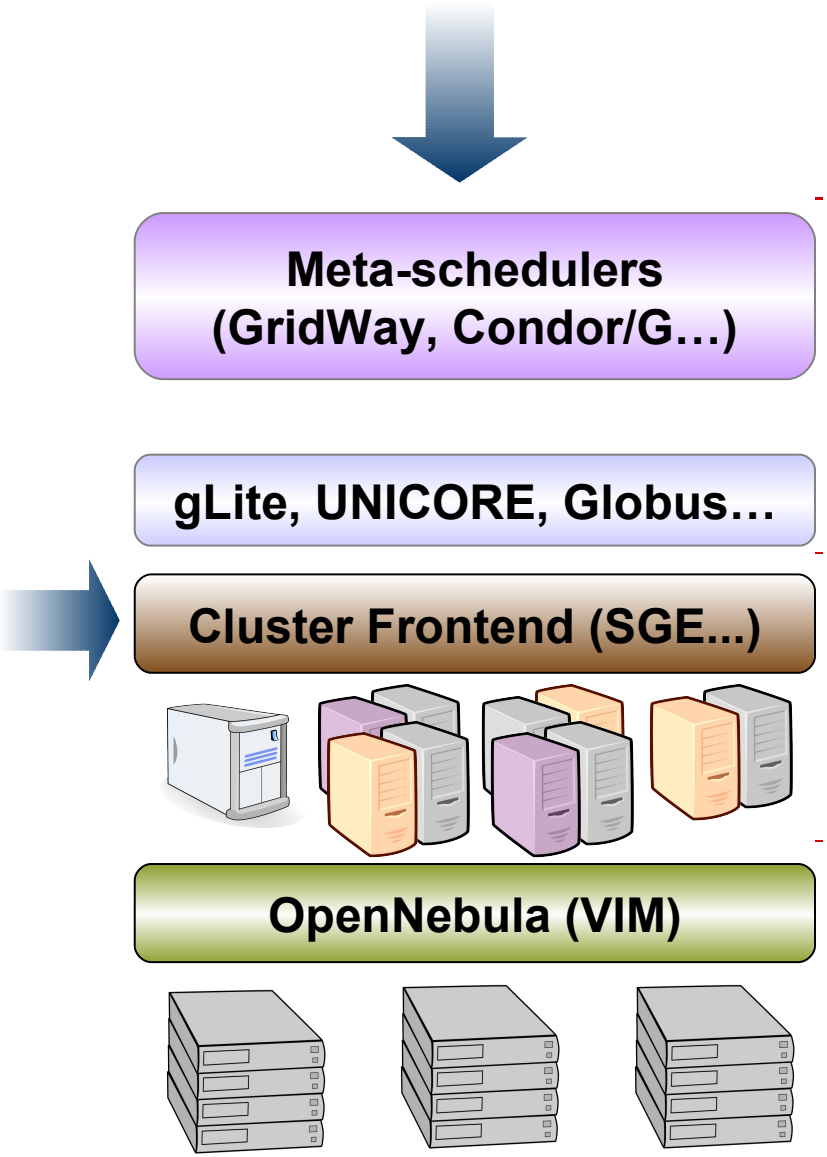

- Unmodified Applications (Grid or local)
- Interfaces preserved (qsub, DRMAA..)

Applications

- Virtual resources are exposed by GM
- Dynamic scheduling
- Fault detection & recovery

Grid Middleware Layer

- WNs register to different queues
- Multiple VO-specific clusters

Computing Service Layer

- Infrastructure consolidation
- Infrastructure partitioning
- Infrastructure adaptation

Infrastructure Layer

A Service to Provide Hardware on Demand (IaaS)

- Cloud systems provide **virtualized resources as a service**
- **On-demand access to infrastructure** (through VMs, and *not jobs*)

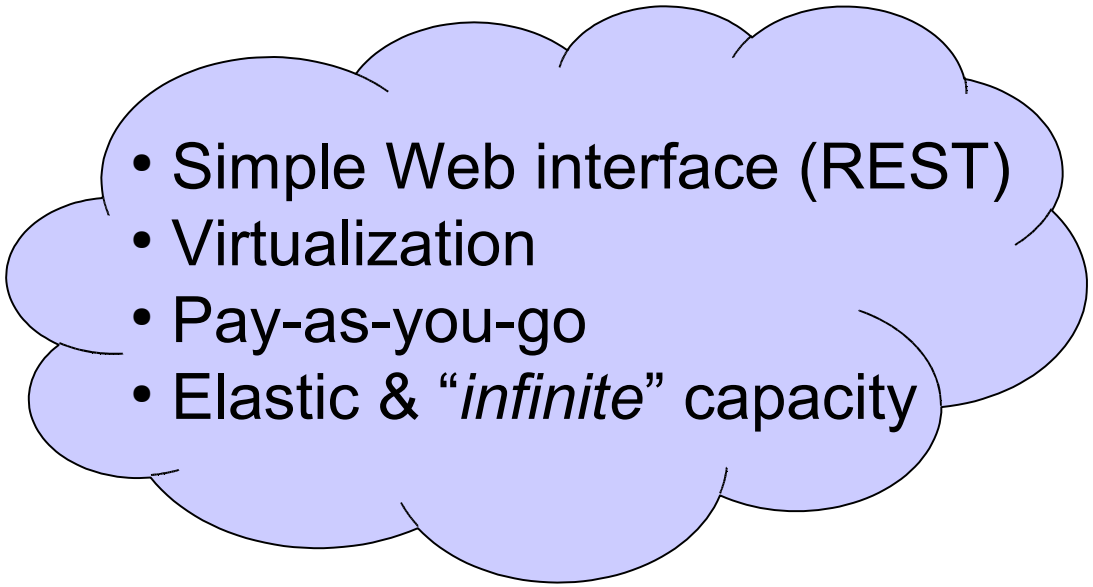

Infrastructure Cloud Services

- Commercial Cloud Providers: Amazon EC2, GoGrid, Elastic Hosts…
- Open Source Cloudn Toolkits: Nimbus, Eucalyptus

Cloud Computing, An Infrastructure View

New provision models for Grids: Virtualization and Clouds

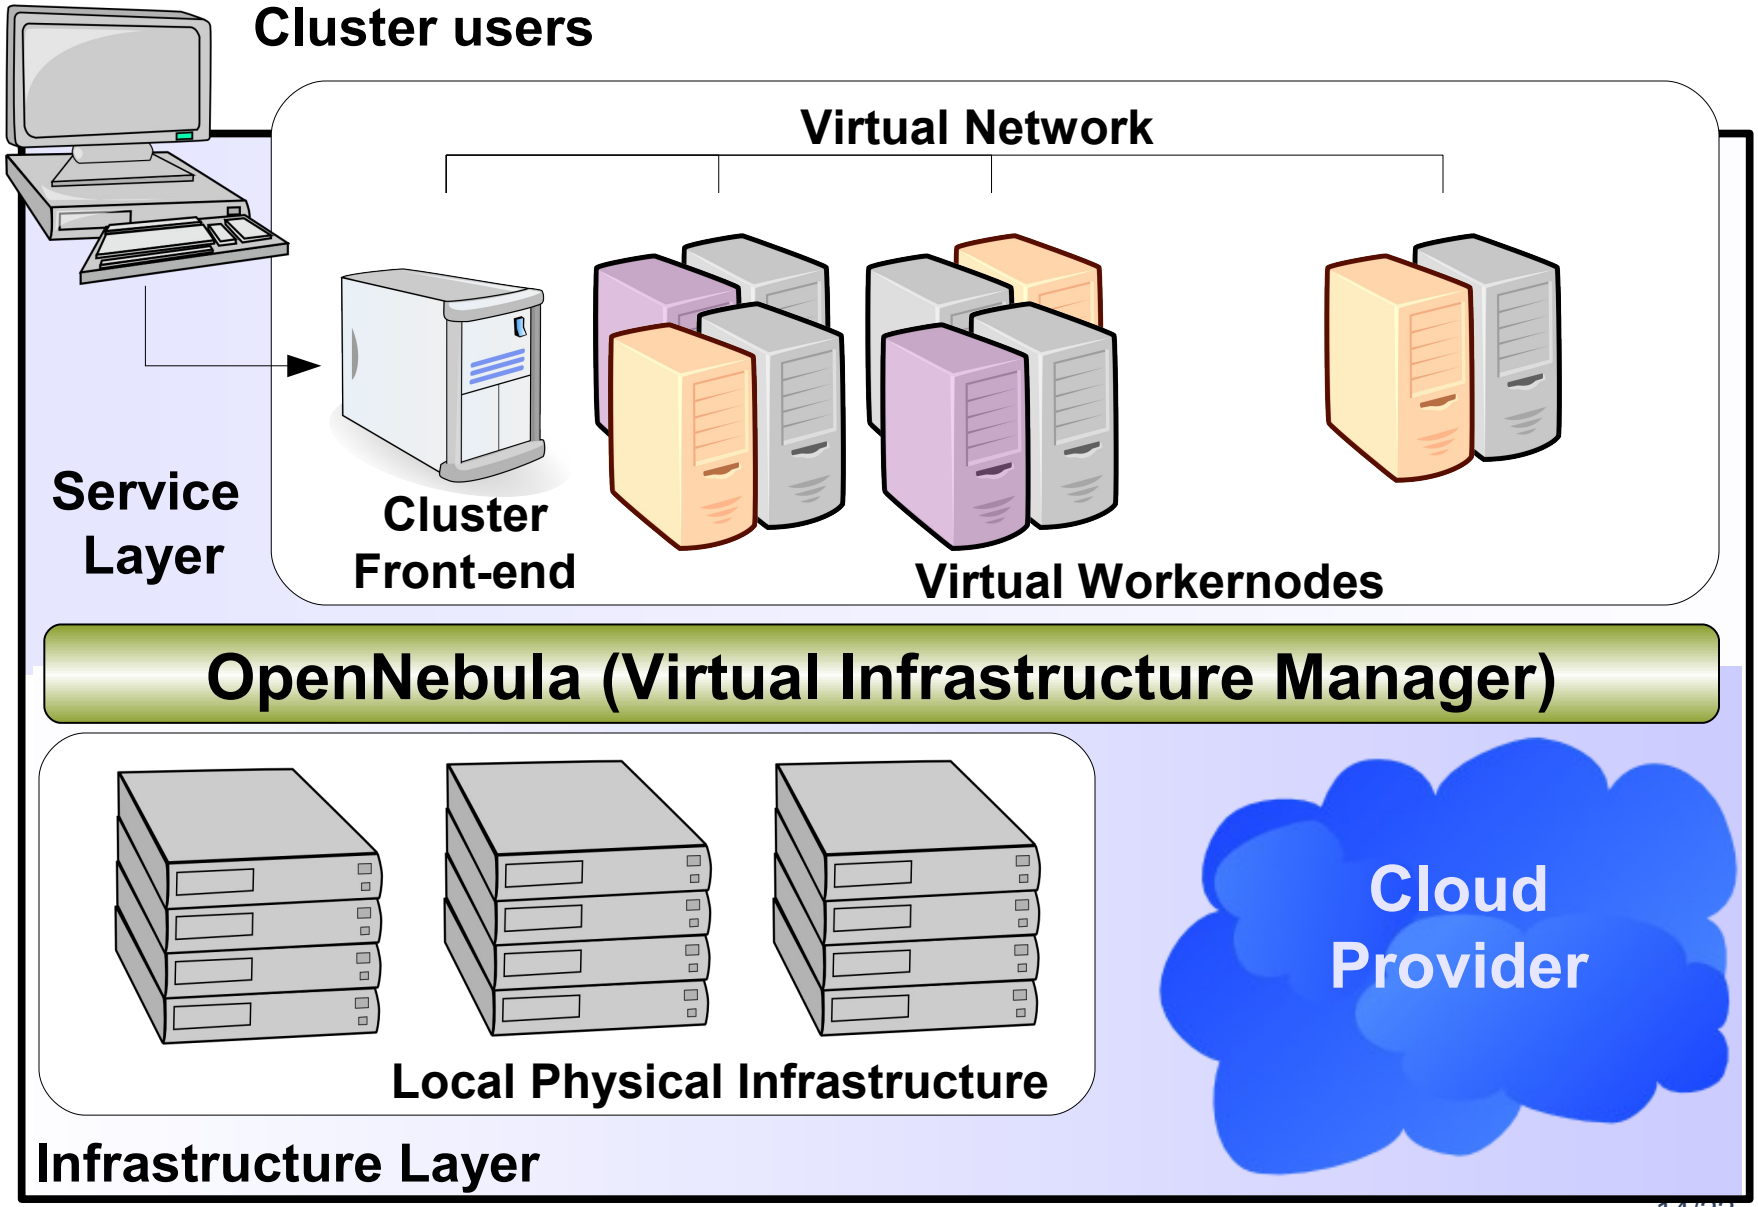

The Infrastructure Quadrant

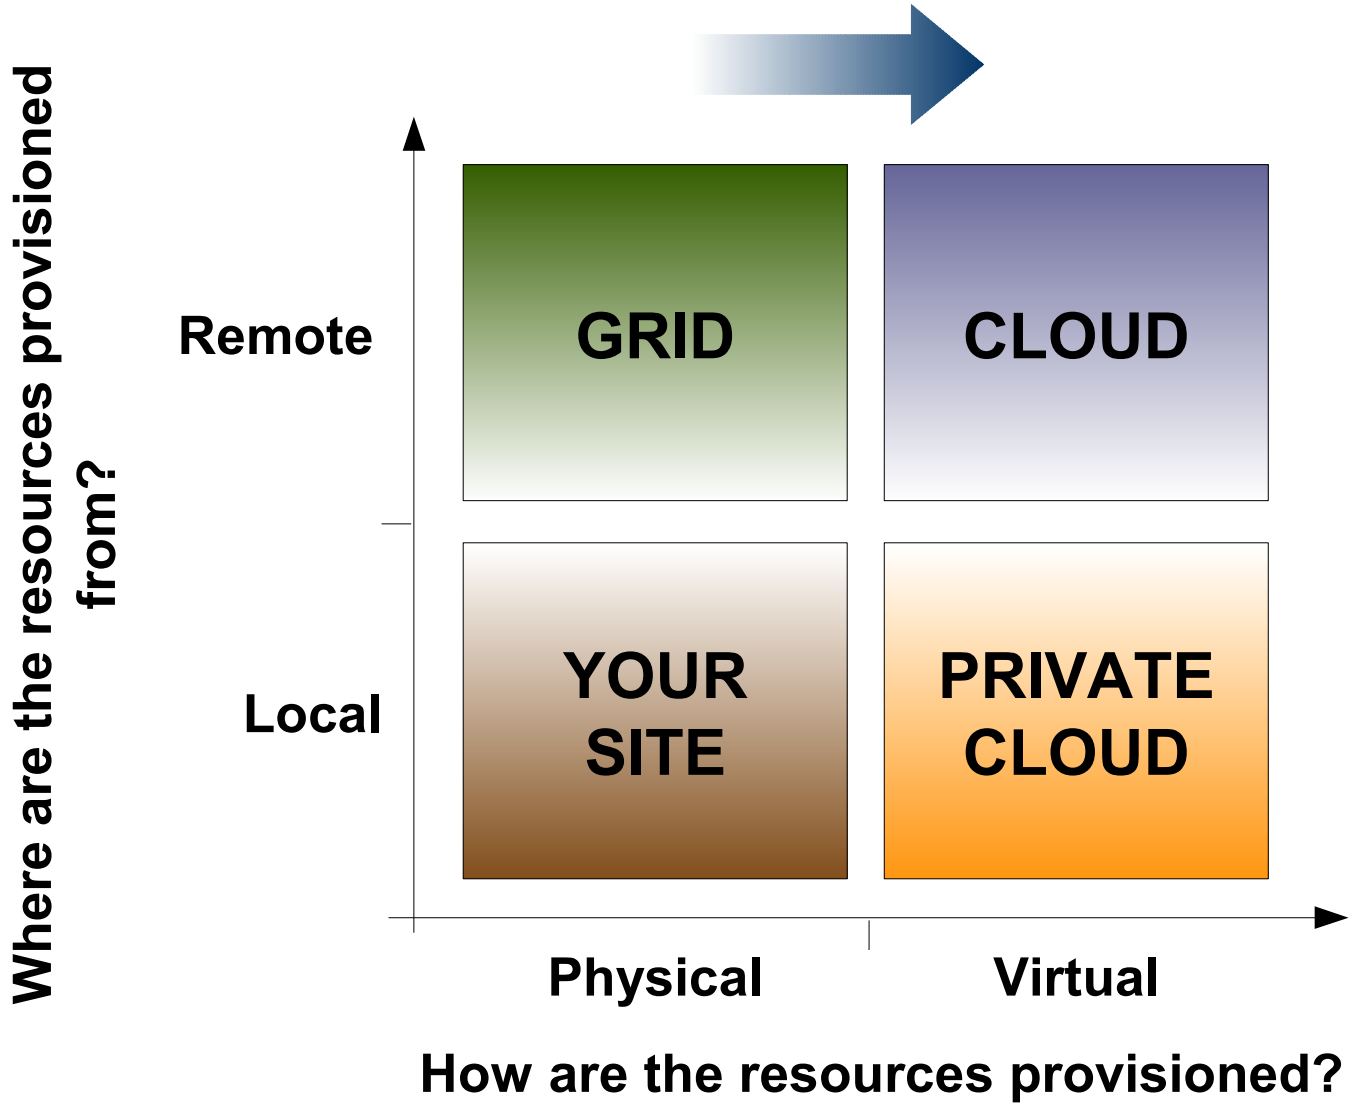

About the Clouds to Support Grids

- Virtualization, cloud, and grid are complementary technologies and will coexist and cooperate at different levels of abstraction
- Virtualization can solve many obstacles for Grid adoption
- Virtualization and cloud do NOT require any modification within service layers (end-user perspective)
- Separation between service and infrastructure layers will allow the application of the utility model to scientific computing *in any form.*

Thanks and More Info...

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

RESERVOIR
www.reservoir-fp7.eu/

The OpenNebula Team

- Ignacio M. Llorente (llorente@dacya.ucm.es)
- Ruben S. Montero (rubensm@dacya.ucm.es)
- Rafel Moreno (rmoreno@dacya.ucm.es)
- Tino Vazquez (tinova@fdi.ucm.es)
- Javier Fontan (jfontan@fdi.ucm.es)

UNIVERSIDAD COMPLUTENSE MADRID

dsa-research.org

THANK YOU FOR YOUR ATTENTION

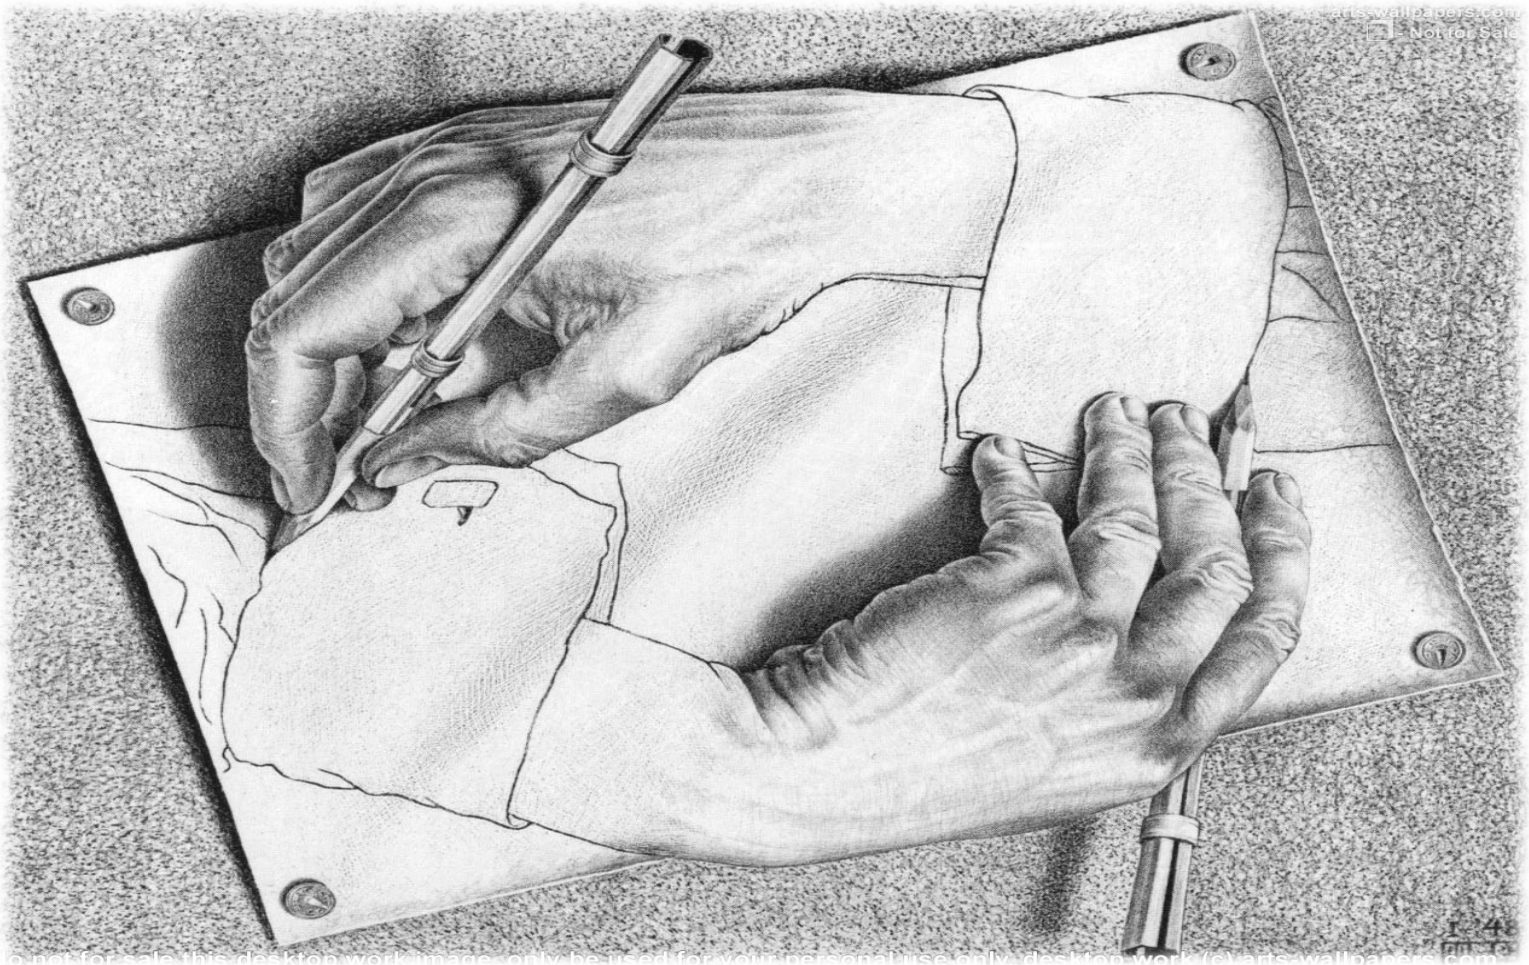

QUESTIONS?