

**OGF-EU: Using IT to reduce Carbon Emissions and  
Delivering the Potential of Energy Efficient Computing  
OGF25, Catania, Italy  
5 March 2009**

**VM Management for Green Data Centres with  
the OpenNebula Virtual Infrastructure Engine**

**Ignacio M. Lorente**

**dsa-research.org**

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





# Objectives

## *VM Management for Green Data Centres with OpenNebula*

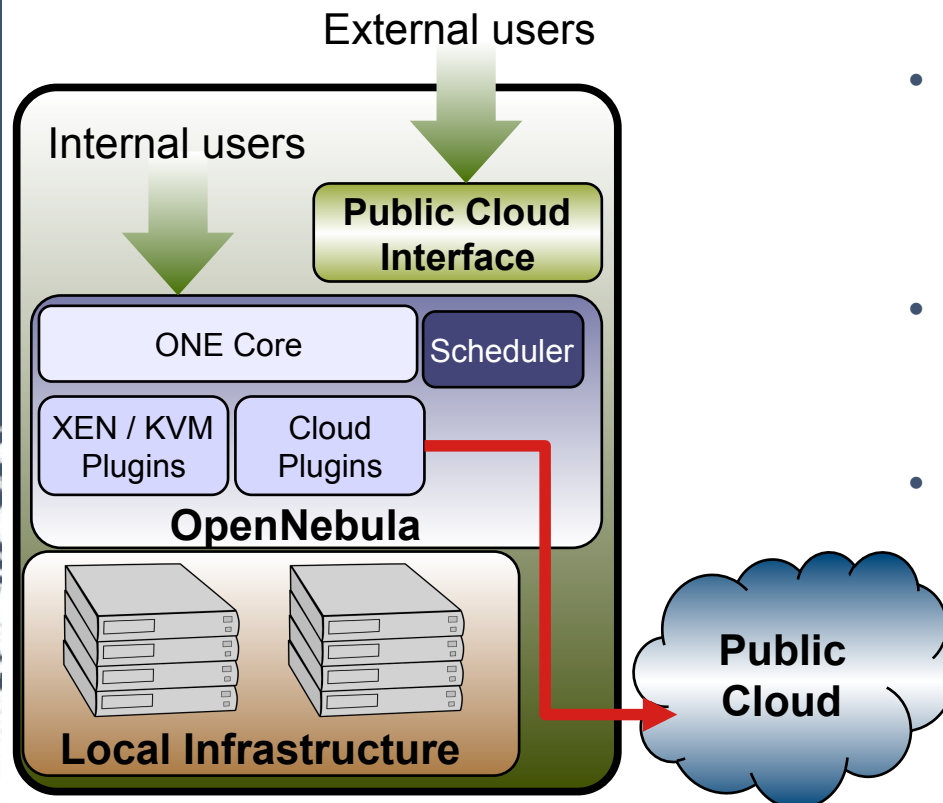
- **Workshop:**
  - Development of a **reference model for the management of energy efficiency** in virtualized distributed environments
- **Presentation:**
  - Introduce the **OpenNebula Virtual Infrastructure Engine**
  - Propose a **model for the dynamic management of VMs in distributed infrastructures**
  - Describe a first prototype of scheduler for **distribution of VM workloads based on energy requirements**

# What is OpenNebula?

*VM Management for Green Data Centres with OpenNebula*

## Extending the Benefits of Virtualization to Clusters

- Dynamic deployment and re-placement of virtual machines on a pool of physical resources
- Transform a rigid distributed physical infrastructure into a **flexible and agile virtual infrastructure**




- **Private Cloud:** Virtualization of cluster or data-center for internal users
- **Backend of Public Cloud:** Internal management of the infrastructure
- **Cloud Interoperation:** On-demand access to public clouds

# Virtual Machine Management Model


VM Management for Green Data Centres with OpenNebula

## Service as Management Entity

- **Service structure**
  - Service components run in VMs
  - Inter-connection relationship
  - Placement constraints
- The VM Manager is **service agnostic**
- Provide **infrastructure context**



## Distributed VM Management Model



# Benefits

## *VM Management for Green Data Centres with OpenNebula*

### System Manager

---

- **Centralized management** of VM workload and distributed infrastructures
- Support for **VM placement policies**: balance of workload, server consolidation...
- **Dynamic resizing** of the infrastructure
- **Dynamic partition** and isolation of clusters
- Support for **heterogeneous workload**
- **Dynamic scaling** of private infrastructure to meet fluctuating demands

### Service Manager

---

- **On-demand provision** of virtual machines

### System Integrators

---

- **Open and flexible** architecture and interfaces, open source software
- **Integration** with any component in the virtualization/cloud ecosystem, such as cloud providers, hypervisors, cloud-like interfaces, virtual image managers, service managers, schedulers...



# Features

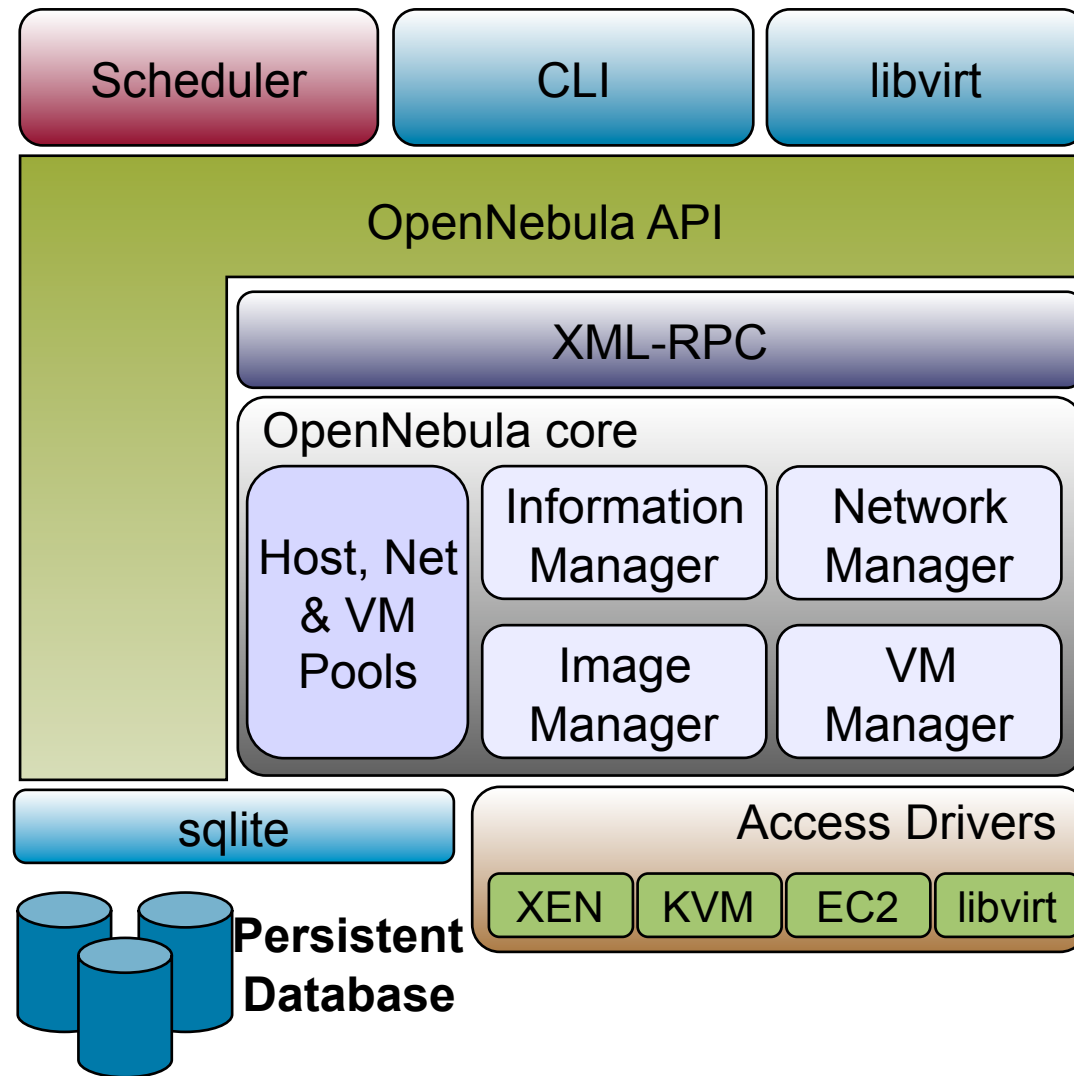
## *VM Management for Green Data Centres with OpenNebula*

Feature	Function
<b>User Interface</b>	<ul style="list-style-type: none"><li>• Unix-like CLI to manage VM life-cycle and physical boxes</li><li>• XML-RPC API and libvirt interface</li></ul>
<b>Scheduler</b>	<ul style="list-style-type: none"><li>• Requirement/rank matchmaker</li><li>• Generic framework to build any scheduler</li></ul>
<b>Virtualization Management</b>	<ul style="list-style-type: none"><li>• Xen, KVM and libvirt connectors</li><li>• Amazon EC2</li></ul>
<b>Image Management</b>	<ul style="list-style-type: none"><li>• General mechanisms to transfer and clone VM images</li></ul>
<b>Network Management</b>	<ul style="list-style-type: none"><li>• Definition of virtual networks to interconnect VMs</li></ul>
<b>Fault Tolerance</b>	<ul style="list-style-type: none"><li>• Persistent database backend to store host and VM information</li></ul>
<b>Scalability</b>	<ul style="list-style-type: none"><li>• Tested in the management of hundreds of VMs</li></ul>
<b>Installation</b>	<ul style="list-style-type: none"><li>• Installation on a UNIX cluster front-end without requiring new services in the remote resources</li><li>• Distributed in Ubuntu 9.04 (Jaunty Jackalope), due to be released in April 2009</li></ul>



# Open and Flexible Architecture

VM Management for Green Data Centres with OpenNebula






# Use Cases

## VM Management for Green Data Centres with OpenNebula


### On-demand Scaling of Computing Clusters

- Elastic execution of a **SGE computing cluster**
- Dynamic growth of the number of worker nodes to meet demands using EC2
- Private network with NIS and NFS
- EC2 worker nodes connect via VPN



### On-demand Scaling of Web Servers

- 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





# Ecosystem

VM Management for Green Data Centres with OpenNebula

## Schedulers

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

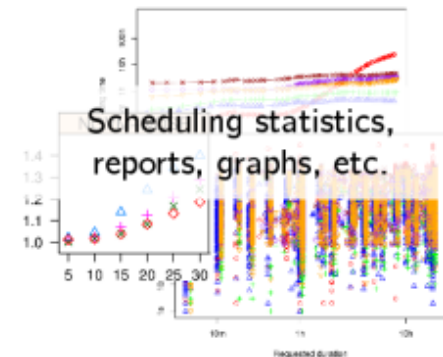


**Lease requests**  
 "I need 10 nodes, each with  
 2 CPUs, 4GB of memory,  
 from 2pm to 4pm"



With  
OpenNebula

Simulation



## Interfaces

- **Libvirt:** Provides an abstraction of a whole cluster of resources as one host, hiding specific hypervisor details.
- **Nimbus:** Can be used as a WSRF or EC2 front-end.

## Plug-Ins

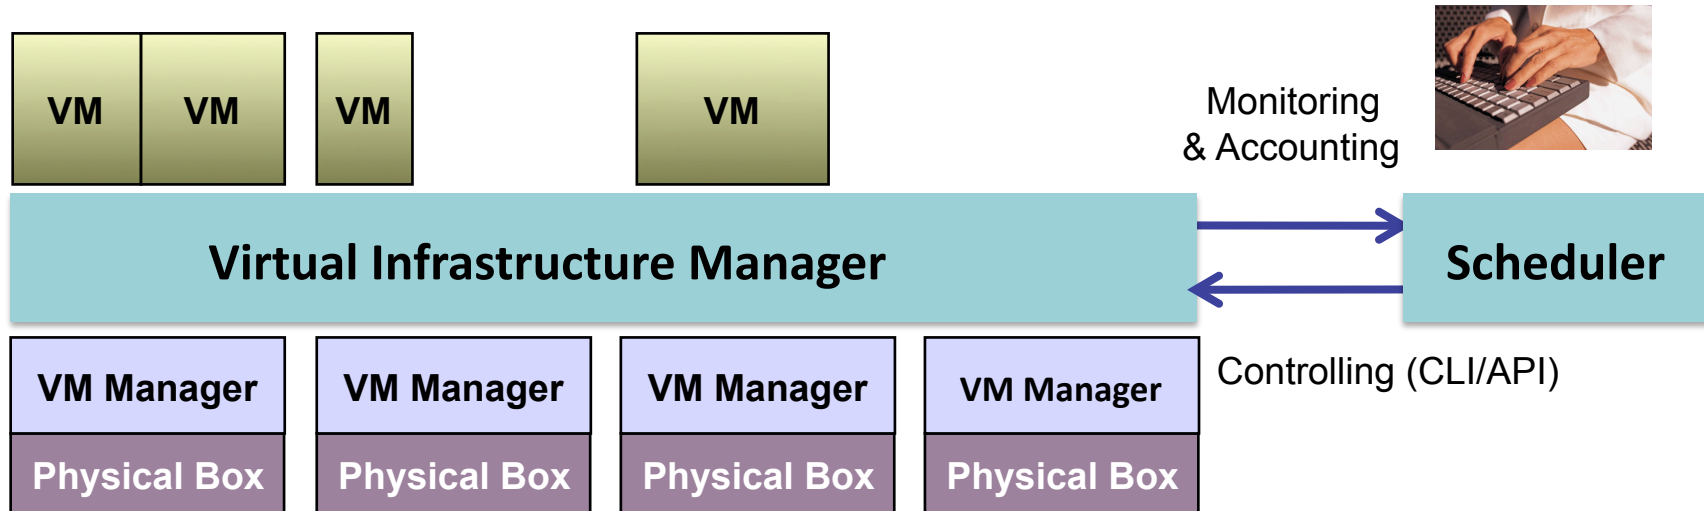
- **ElasticHosts:** Enables the dynamically increase capacity of your virtualized infrastructure to meet fluctuating peak demands using a cloud provider.

# Model for VM Management in Distributed Infrastructures

VM Management for Green Data Centres with OpenNebula

## Energy Policy Enactment

## Energy Policy Definition



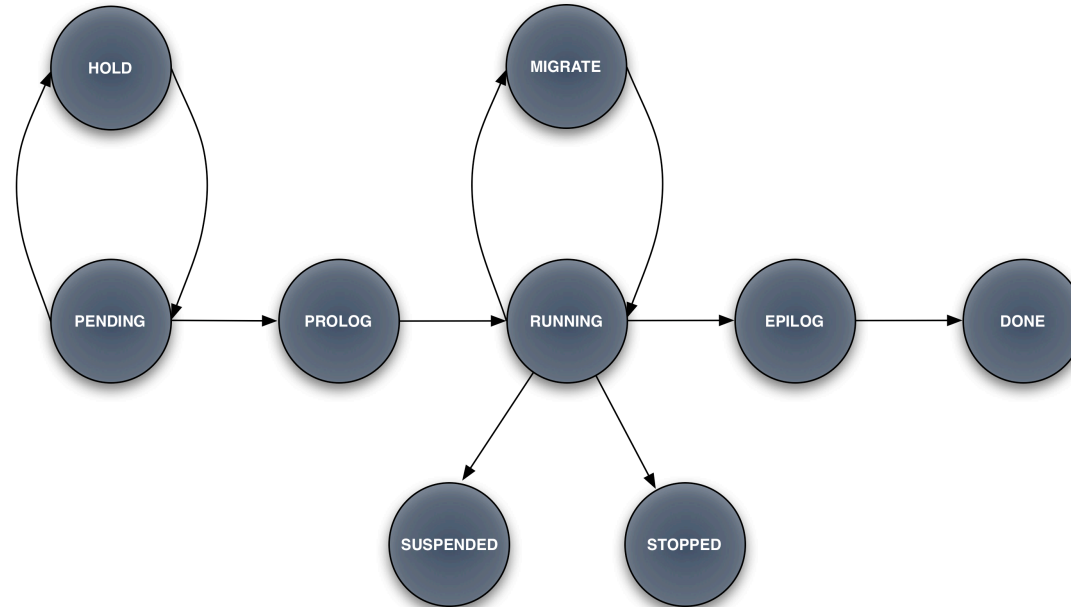
### Centralized management of VMs and resources

- VM life-cycle management
- VM image management
- Virtual network management
- Fault tolerance

### Automation of VM placement:

- Balance of workload
- Server consolidation
- Placement constraints and affinity
- Advance reservation of capacity
- SLA commitment

### VM Life-cycle and Monitoring



Attribute	Description
ID	VM identifier
NAME	Name of the VM
STAT	Status
CPU	CPU percentage used by the VM
MEM	Memory used by the VM
HOSTNAME	Host where the VM is running
TIME	Time since submission

# Monitoring and Accounting

*VM Management for Green Data Centres with OpenNebula*

## Resource Monitoring

Attribute	Description
HID	Host ID
NAME	Host name
RVM	Number of running VMs
TCPU	Total CPU (percentage)
FCPU	Free CPU (percentage)
ACPU	Available CPU (not allocated by VMs)
TMEM	Total memory
FMEM	Free memory
STAT	Host status

## Flexible Architecture

- Easy addition of new VM and resource metrics (energy efficiency metrics)
- Experimenting with Advanced Configuration and Power Interface (ACPI) for power monitoring

## Accounting

- Statistics of host usage, VM image transfer times....

# Controlling

## VM Management for Green Data Centres with OpenNebula

### VM Controlling

Action	Description
create	Submits a new virtual machine, adding it to the VM pool
deploy	Starts a previously submitted VM on a specific host
shutdown	Shutdown an already deployed VM
livemigrate	Migrates a running VM to another host without downtime
migrate	Saves a running VM and starts it again in the specified host
stop	Stops a running VM
Release/hold	Releases/hold a VM from/to hold state
Suspend/resume	Saves/resume a VM

### Resource Controlling

Action	Description
Create/delete	Adds/removes a machine to/from the pool
Enable/disable	Enables/disables host



# Resource Schedulers

*VM Management for Green Data Centres with OpenNebula*

## Requirement/rank Matchmaker (default)

- First placement of pending VMs to resources meeting the “requirements” and sorted using a “rank” expression

## Haizea Lease Manager (University of Chicago)



- Advance reservation of capacity and queuing of best effort requests

## RESERVOIR Policy Engine (IBM Haifa/Elsa Datamat)



- Policy-driven probabilistic admission control and dynamic placement optimization to satisfy site level management policies

## VM Consolidation Scheduler

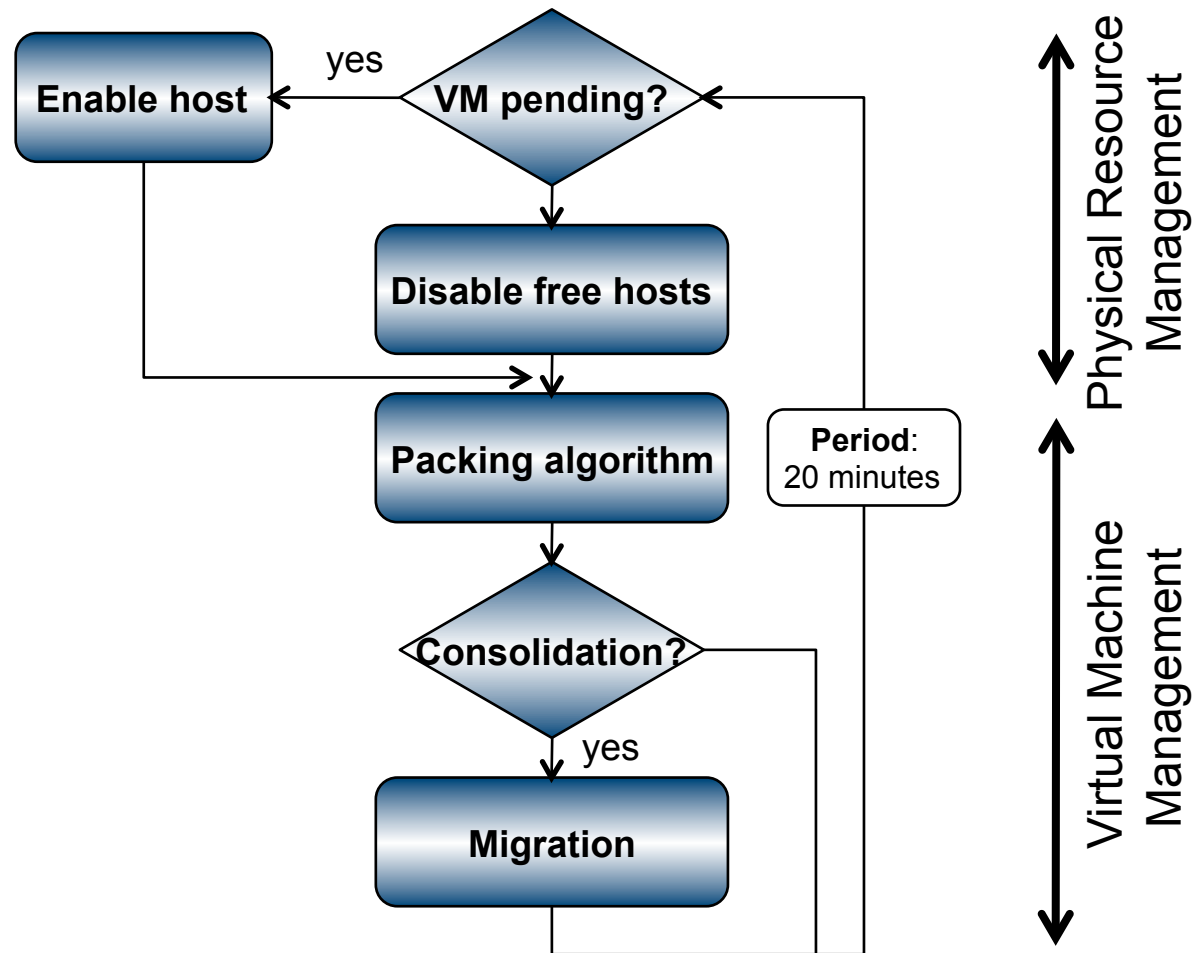
- Periodic re-placement of VMs for server consolidation and suspension/resume of physical resources
- *Experimental* for supporting research on distribution of VM workload based on energy requirements and policies

# VM Consolidation Scheduler

VM Management for Green Data Centres with OpenNebula

## Control Flow

- Complements the work done by the requirement/rank matchmaker for first deployment of VMs in pending state

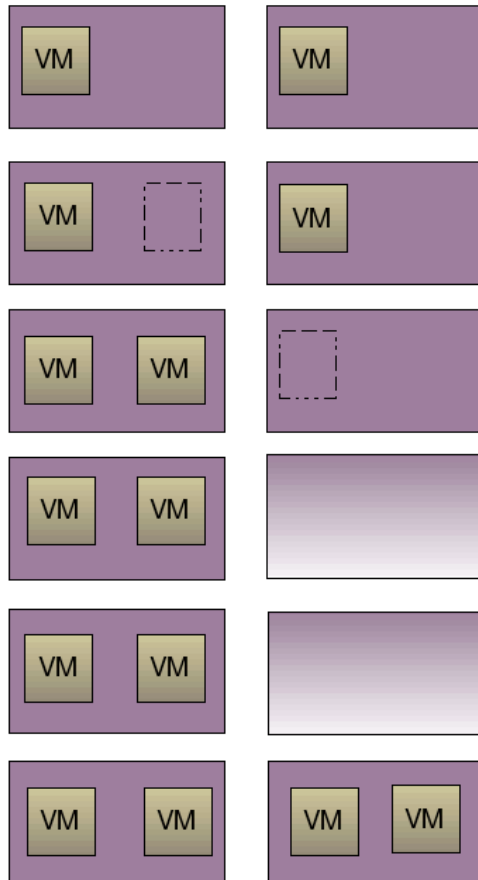




# VM Consolidation Scheduler

VM Management for Green Data Centres with OpenNebula

## A Simple Example



1) Initial situation

2)

**Exec. 1:** Packing algorithm in the VM Consolidation Scheduler requests a migration

3)

**Exec. 2:** VM Consolidation Scheduler “disables” a physical resource

4)

**Exec. 3:** VM consolidation scheduler “enables” a physical resource because there are pending VMs



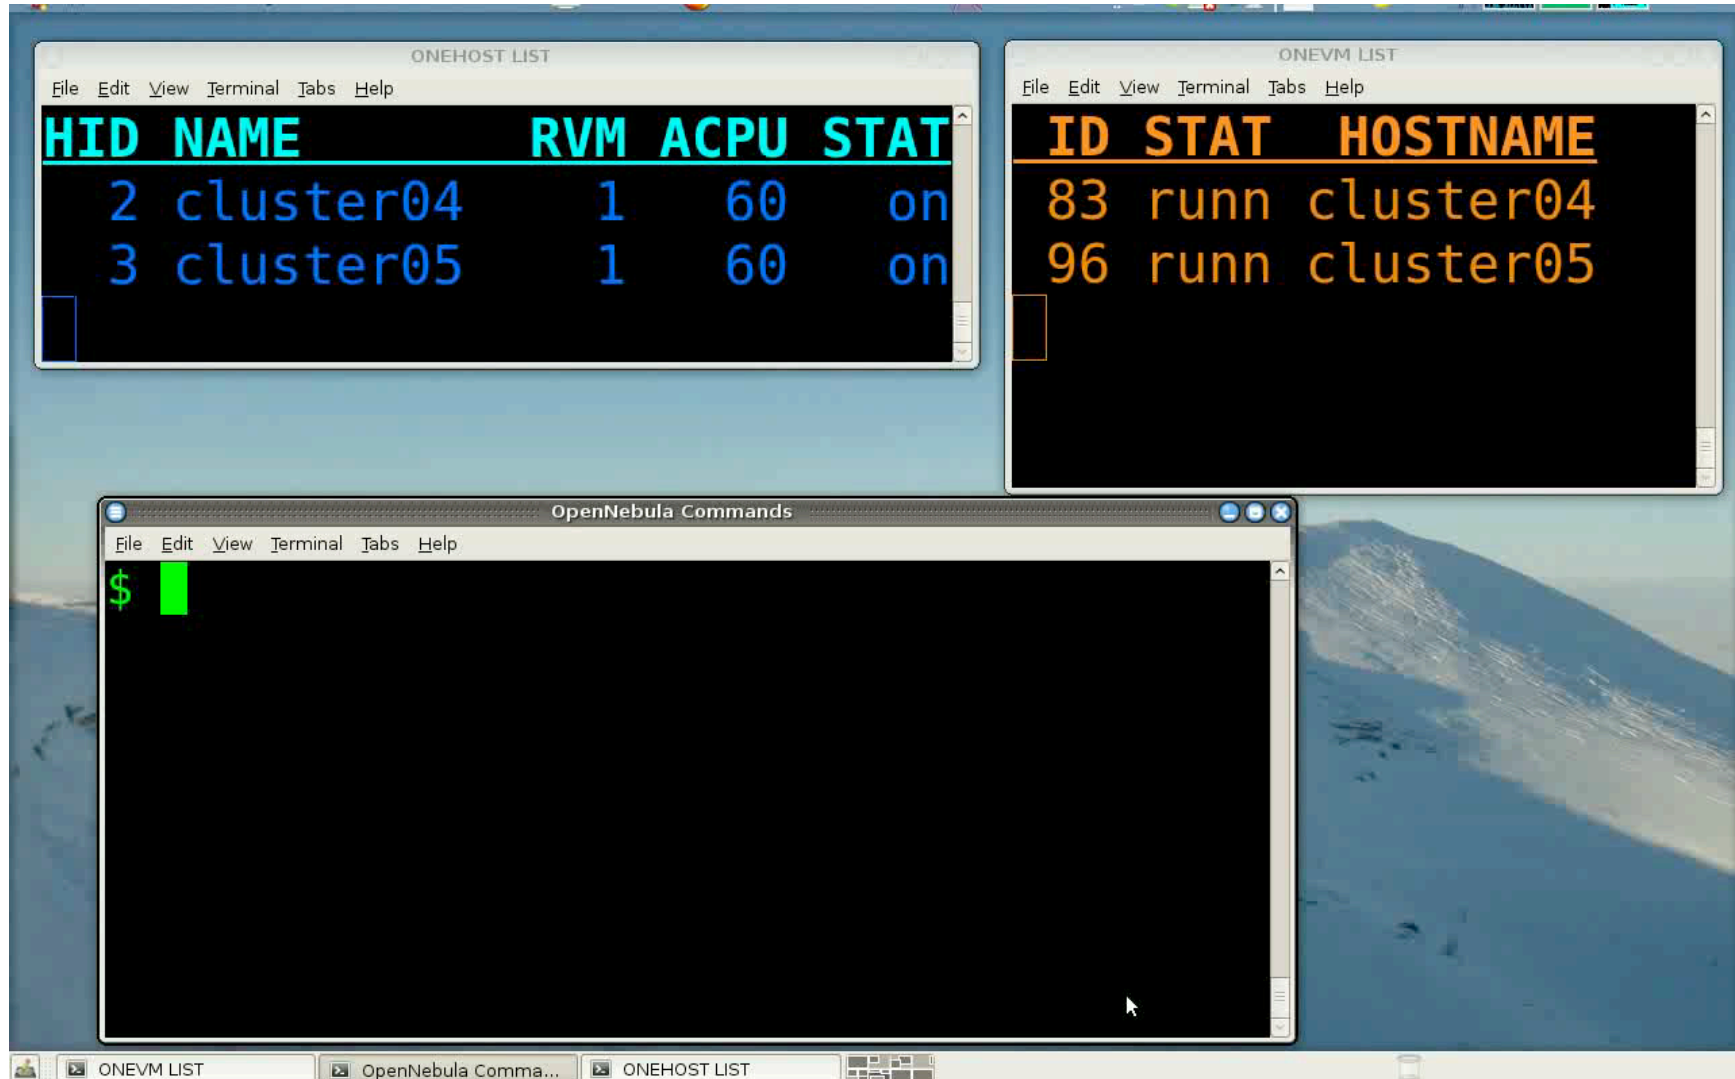
5)

6) The requirement/rank matchmaker schedules the pending VMs

# VM Consolidation Scheduler

VM Management for Green Data Centres with OpenNebula

## A Simple Example



The screenshot displays the OpenNebula web interface with three main components:

- ONEHOST LIST:** A table showing the status of hosts.
- ONEVM LIST:** A table showing the status of virtual machines.
- OpenNebula Commands:** A terminal window with a shell prompt.

HID	NAME	RVM	ACPU	STAT
2	cluster04	1	60	on
3	cluster05	1	60	on

ID	STAT	HOSTNAME
83	runn	cluster04
96	runn	cluster05

The terminal window shows a green prompt character '\$' followed by a red cursor bar.



# VM Consolidation Scheduler

*VM Management for Green Data Centres with OpenNebula*

## Future Work: Energy Policy-driven Scheduler

### Monitoring

- Power and heat metrics (ACPI specification)
- Past history (trends) of resources and workload

### Enactment of policies

- Management of power-state of the resources (ACPI specification)

### Policies

- Support for “generic energy policies”
- Selection of the VMs to migrate and the target host for migration
- ...



# Conclusions

## *VM Management for Green Data Centres with OpenNebula*

- **Workshop:**
  - Development of a **reference model for the management of energy efficiency** in virtualized distributed environments
- **Presentation:**
  - Introduce the **OpenNebula Virtual Infrastructure Engine**
  - Propose a **model for the dynamic management of VMs in distributed infrastructures**
  - Describe a first prototype of scheduler for **distribution of VM workloads based on energy requirements**



**THANK YOU FOR YOUR ATTENTION!!!**  
More info, downloads, mailing lists at  
[www.OpenNebula.org](http://www.OpenNebula.org)

**Real demo at booth number 4**

### Acknowledgements

---

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