

**NASA Ames**

**NASA Advanced Supercomputing (NAS) Division**  
California, May 24th, 2012

**Cloud Computing Architecture with  
OpenNebula  
HPC Cloud Use Cases**

**Ignacio M. Llorente**

Project Director

**OpenNebula.org**


---

*Cloud Computing Architecture with OpenNebula - HPC Cloud Use Cases*


This presentation is about:

- Cloud case studies in HPC
- Private cloud computing using OpenNebula
- The anatomy of a Private Cloud
- OpenNebula innovative features for HPC

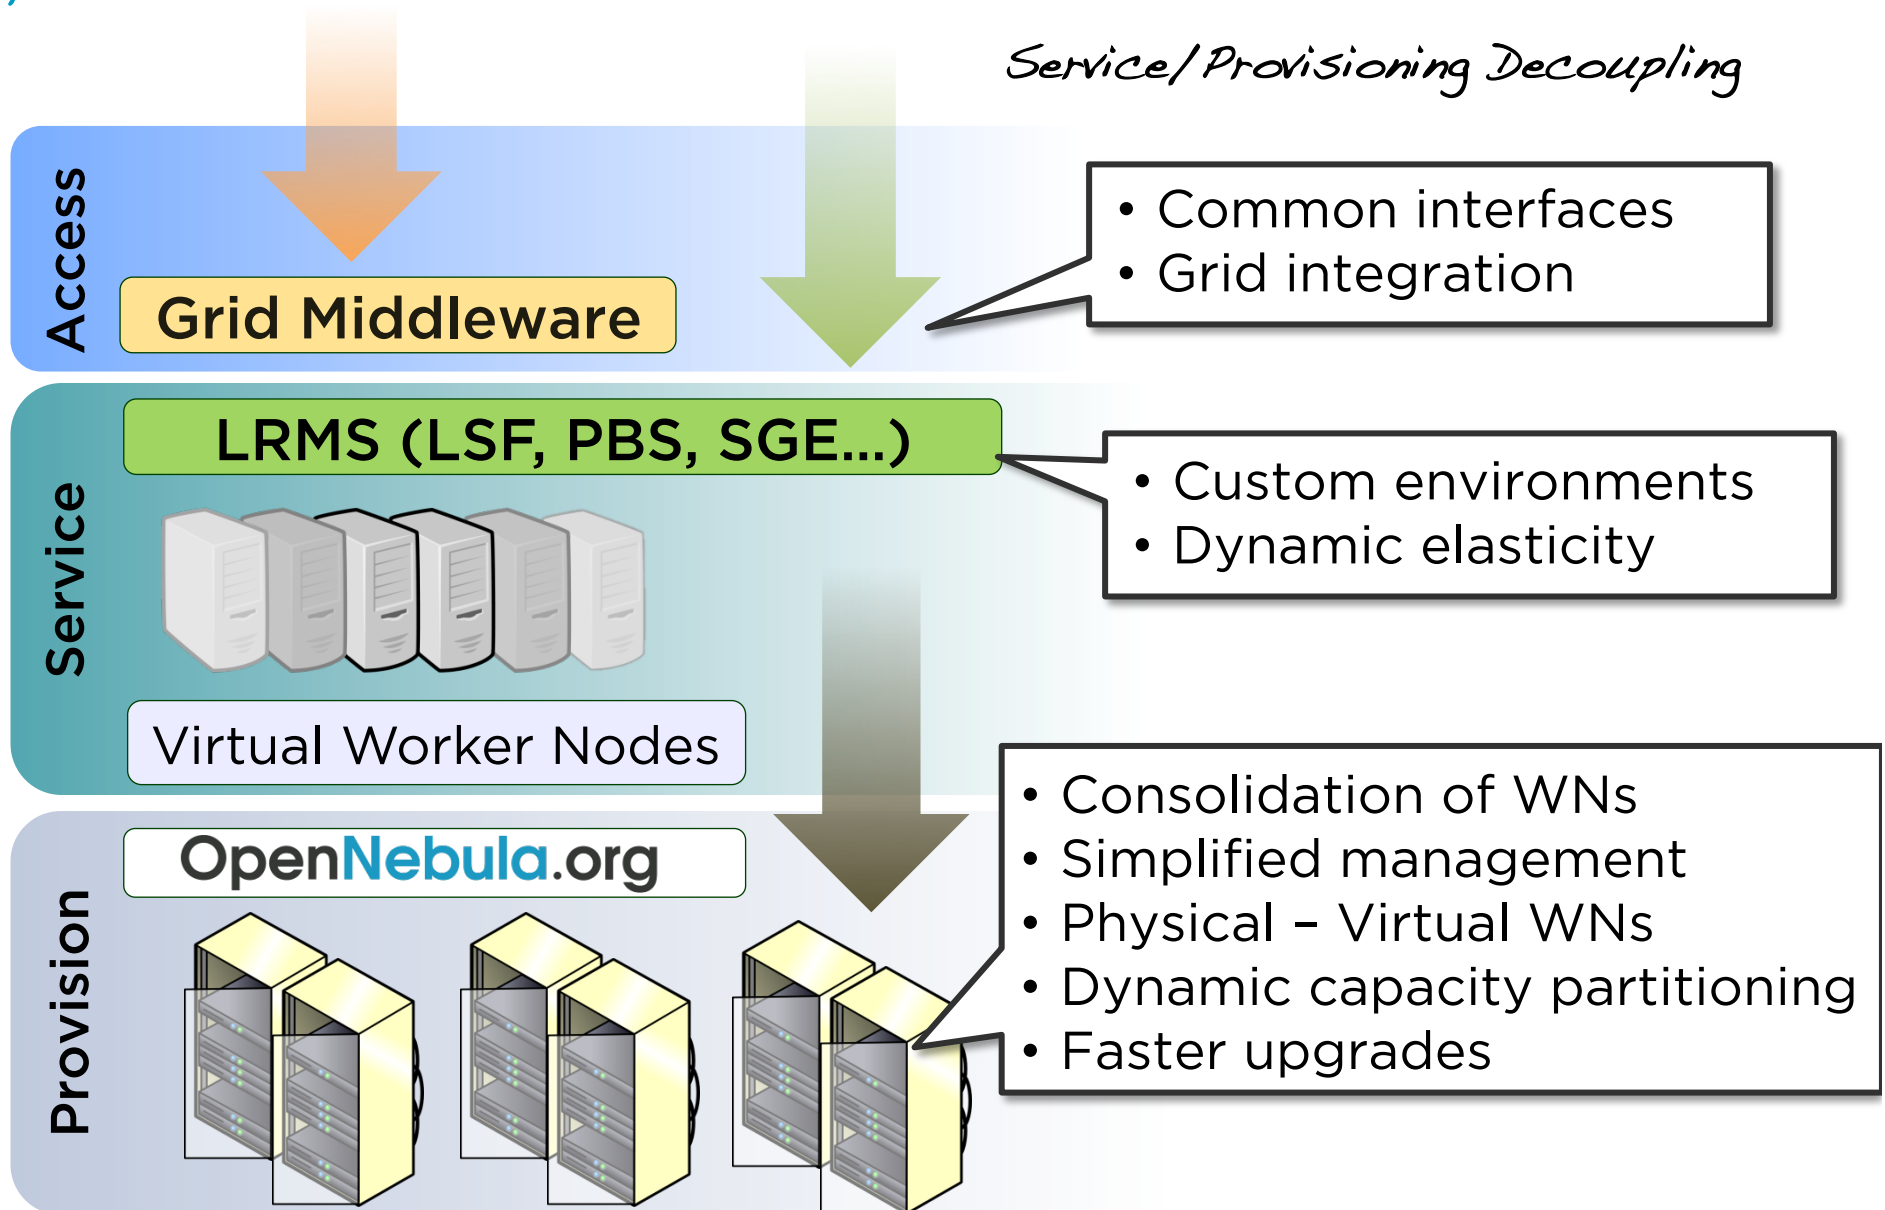
*Open Cloud Solution for Building and Managing Virtualized Data Centers*



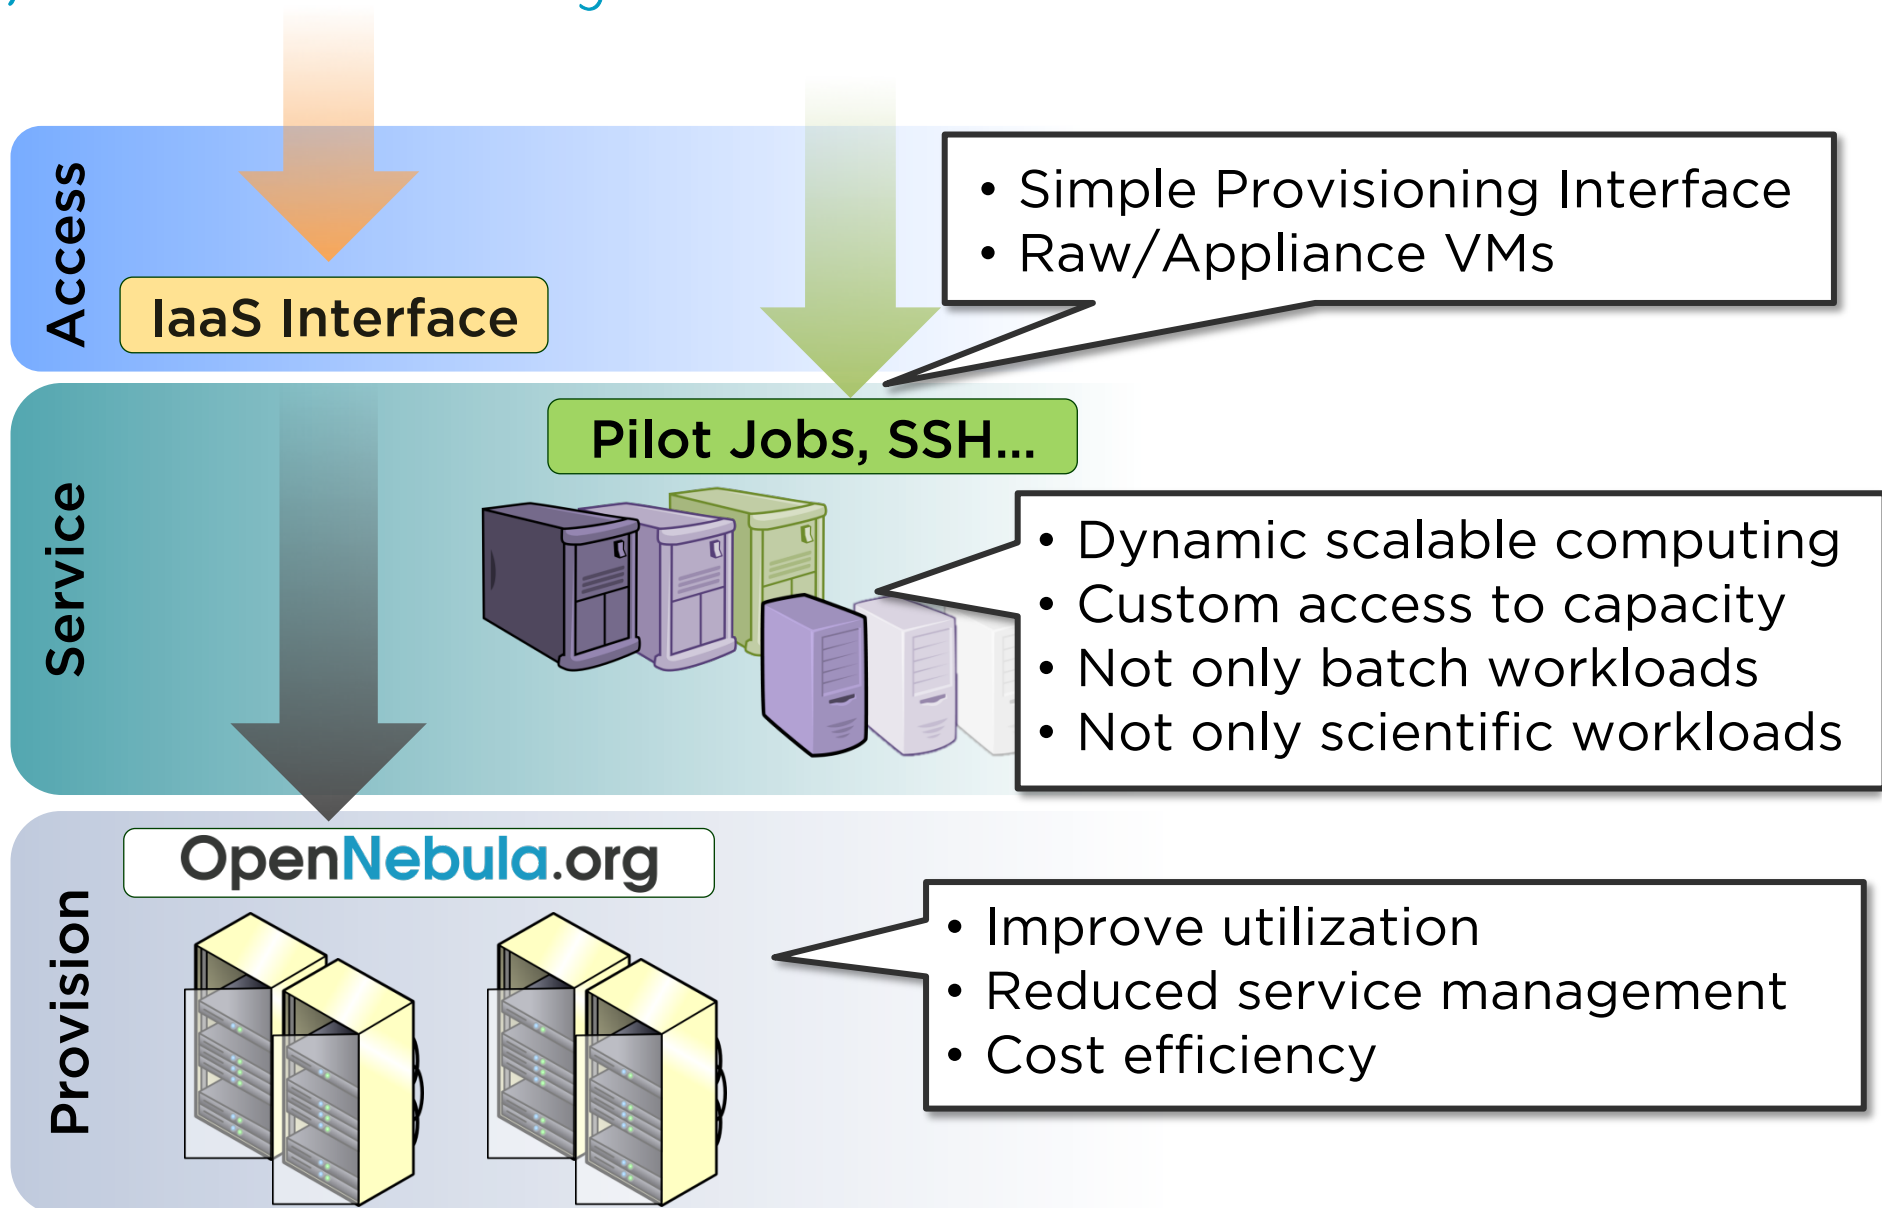
*The Pre-cloud Era*



## OpenNebula as an Infrastructure Tool



## *OpenNebula as an Provisioning Tool*



Examples: CERN's Ixcloud

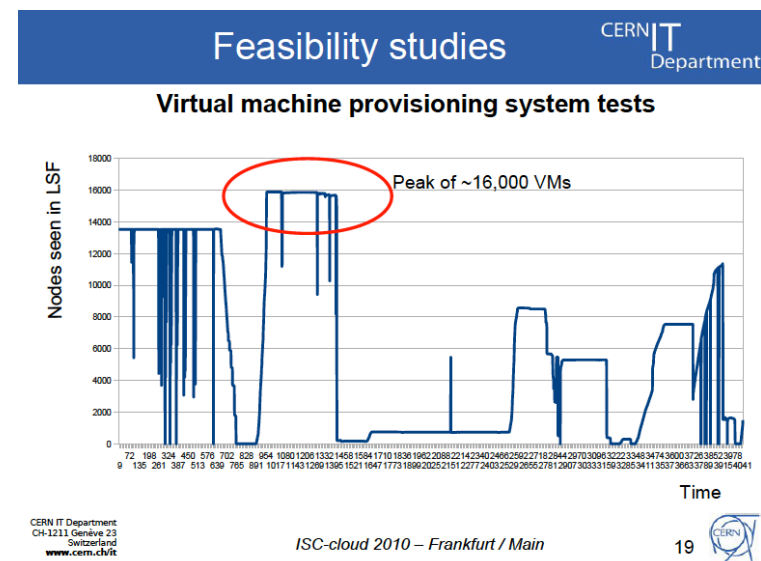
<http://blog.opennebula.org/?p=620>

## Goal

- Virtualized HTC Batch Nodes: limited life-time, dynamic capacity
- EC2 Query: CernVM

## Deployment Notes

- Custom network integration MAC/IP pinning
- Fast image distribution & boot: BitTorrent + LVM snapshots
- OpenNebula tests up to 20,000 VMs (**great feedback!**)



Examples: FermiCloud

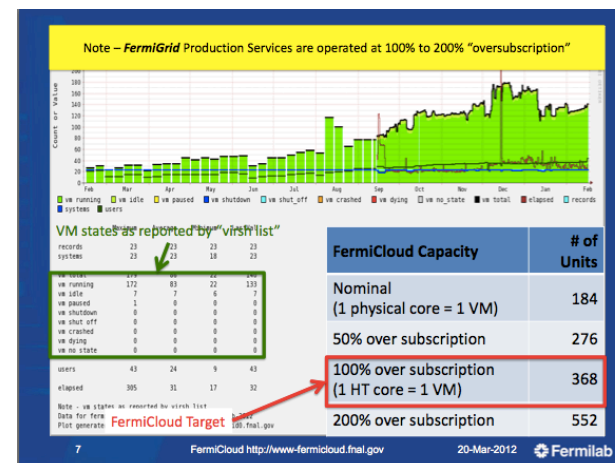
<http://www-fermicloud.fnal.gov/>

## Goal

- Scientific stakeholders get access to on-demand VMs
- Developers & integrators of new Grid applications
- MPI and legacy applications

## Deployment Notes

- VMs access Fermilab Networking and Storage Services
- OpenNebula + X509 support (**contributed back!**)
- Other areas: HA, Batch queues look-ahead, cluster on-demand...





Examples: SARA

<https://www.cloud.sara.nl/>

## Goal

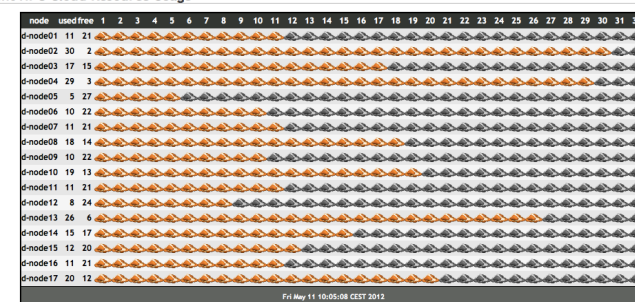
- HPC clusters on-demand: Self service, elastic and fully configurable HPC systems
- Offer production infrastructure for several engineering and scientific communities: Bioinformatics, ecology...

## Deployment Notes

- Low latency network for HPC
- Extensions in accounting, GUI, firewalling... **(contributed back!)**



Current HPC-Cloud Resource Usage



A total of 263 cores out of a maximum of 544 cpu cores are in use at this time. (48%).



The largest VM that can currently be deployed can use at most 27 cores.



The number of Virtual Machines currently running.

## One of Our Main User Communities

### Supercomputing Centers




### Research Centers



### Distributed Computing Infrastructures



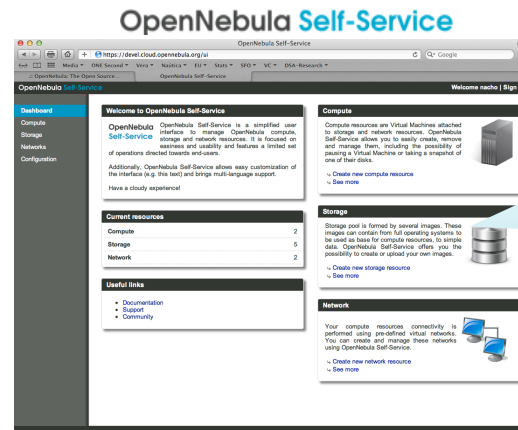
## *Different Perspectives of the Cloud*





## How Can I Create and Manage my Virtual Infrastructure?

... standards (*de facto* and *de jure*) Cloud APIs to leverage existing ecosystems and ensure portability across providers, and **self-service portal** ...

- ! Use AWS existing tools like ElasticFox
- ! Execute complete computing cluster
- ! Manage persistent scientific data
- ! Catalog with pre-defined environments
- ! Sharing between users or groups




• Manage images, networks and instances



*What are the Main Components to Build a Cloud Infrastructure?*

## Instance Networks

- Guests
- Public and private networks



## Front-end


- Authentication
- Authorization
- ACLs, roles, groups...
- Accounting
- Logging
- Resource quotas

## Service Networks

- Monitoring, control...
- Live migration...
- Storage access...



- ! Fair share of resources
- ! Place VM close to data



## Hosts

- Multiple hypervisors
- Up to 500 hosts
- Automatic failover and HA
- Automatic resource allocation



## Datastores


- VM image storage
- Multiple datastores
- Heterogeneous configurations
- Shared or non-shared



## Broad Commodity and Enterprise Platform Support

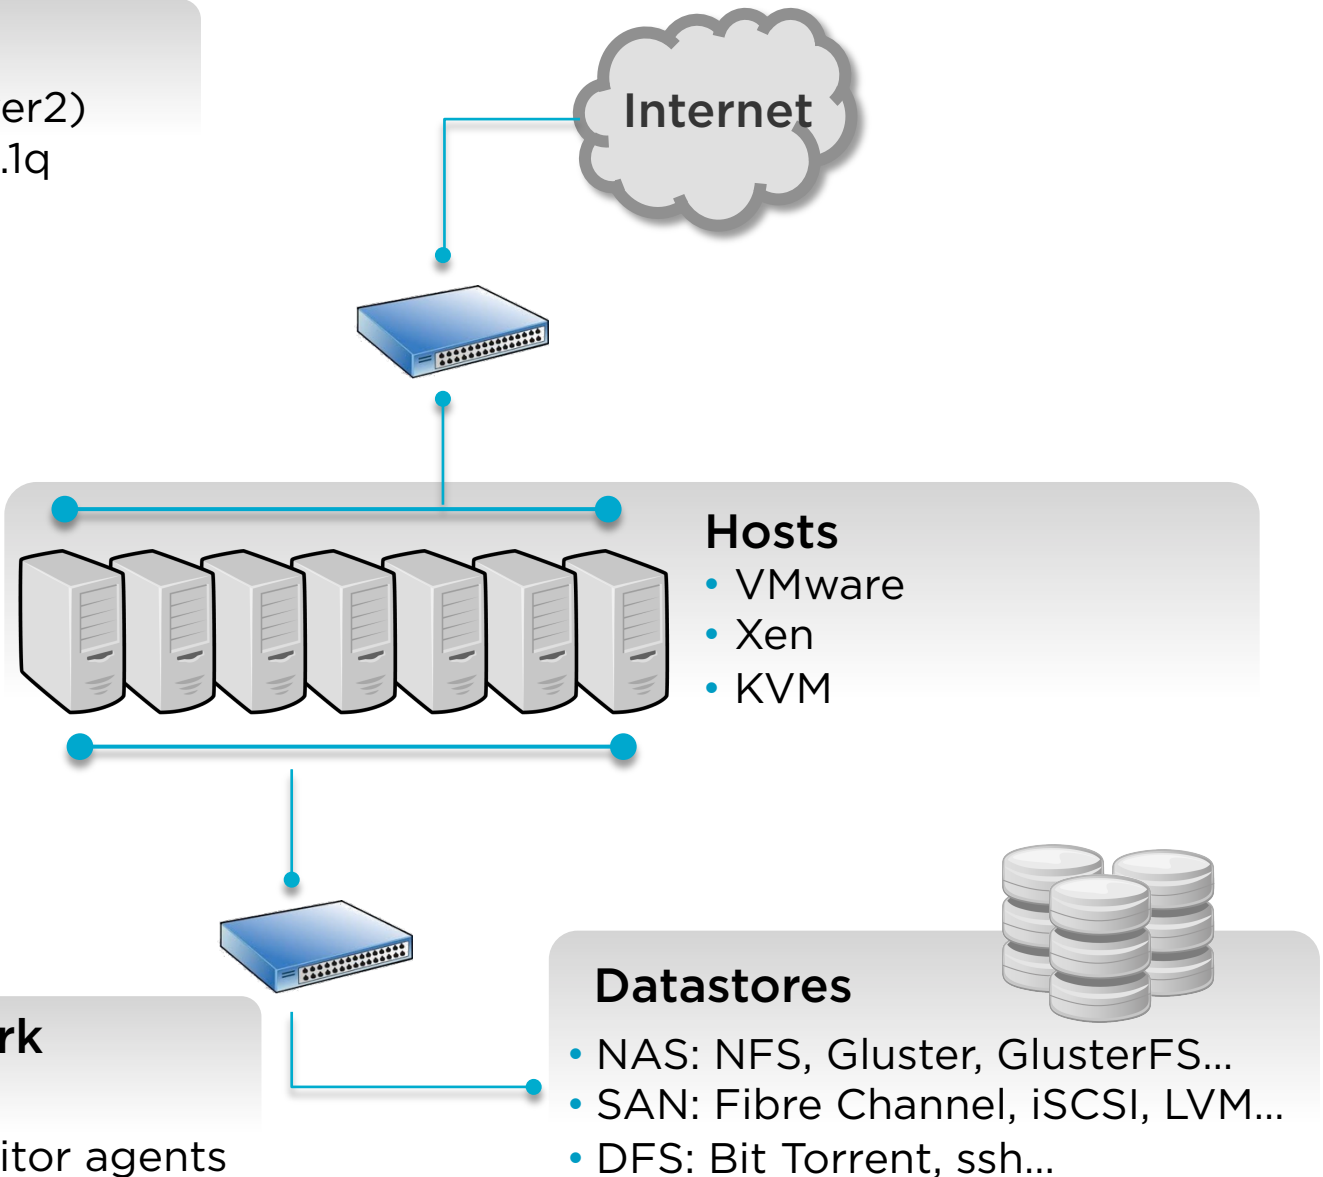
### Instance Network

- VLAN per user (layer2)
- Open vSwitch, 802.1q
- Ebttables



### Front-end

- X509, LDAP, ssh keys
- ACLs, roles, groups...



### Hosts

- VMware
- Xen
- KVM

### Service Network

- Ganglia/Nagios
- Additional monitor agents


### Datastores

- NAS: NFS, Gluster, GlusterFS...
- SAN: Fibre Channel, iSCSI, LVM...
- DFS: Bit Torrent, ssh...

## Clustering the Physical Resources

### Clusters

- **Pools of hosts** that share datastores and networks
- Group different configurations for **performance** and **security**




- ! Different Hw for different workloads profiles (HPC vs HTC)
- ! Isolation and security levels

### Multiple Datastores per Cluster

- **Balance I/O operations** between storage servers
- Define different **SLA policies** (e.g. backup) and **performance features** for different VM types or users



## Centralized Management of Multiple OpenNebula Instances (Zones)



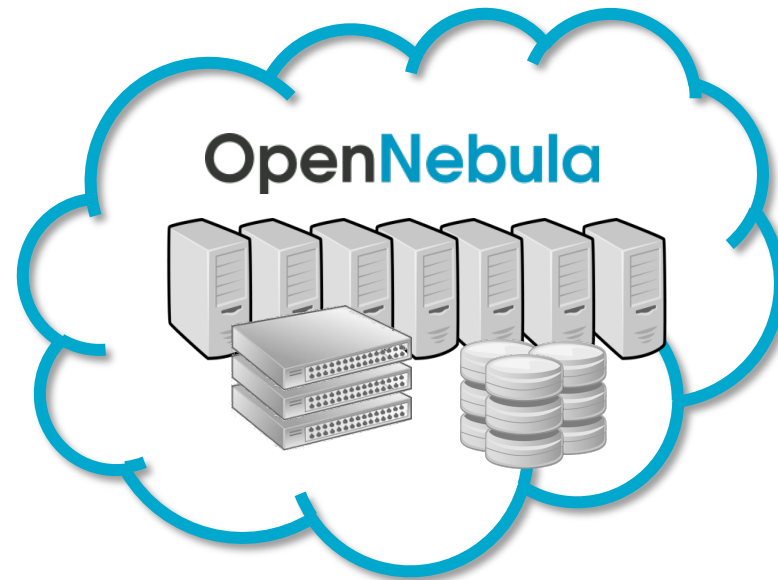

Cloud Consumer



### oZones Server

- Portal
- Cloud API (EC2, OCCl)
- Global AuthN

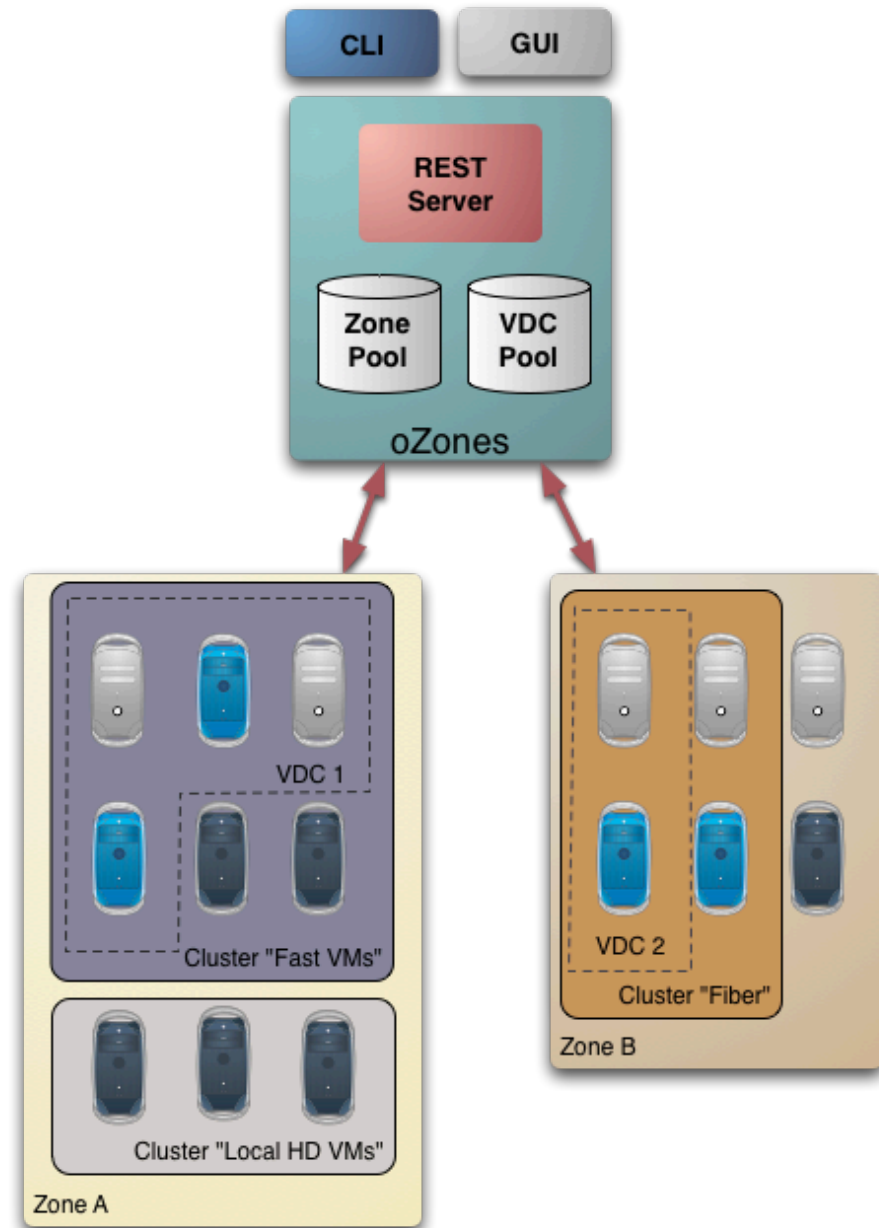

- ! Full Isolation of performance-sensitive applications
- ! Isolation and security levels
- ! Multi-tier
- ! Scalability
- ! Multiple-site





## On-demand Provision of Virtual Data Centers

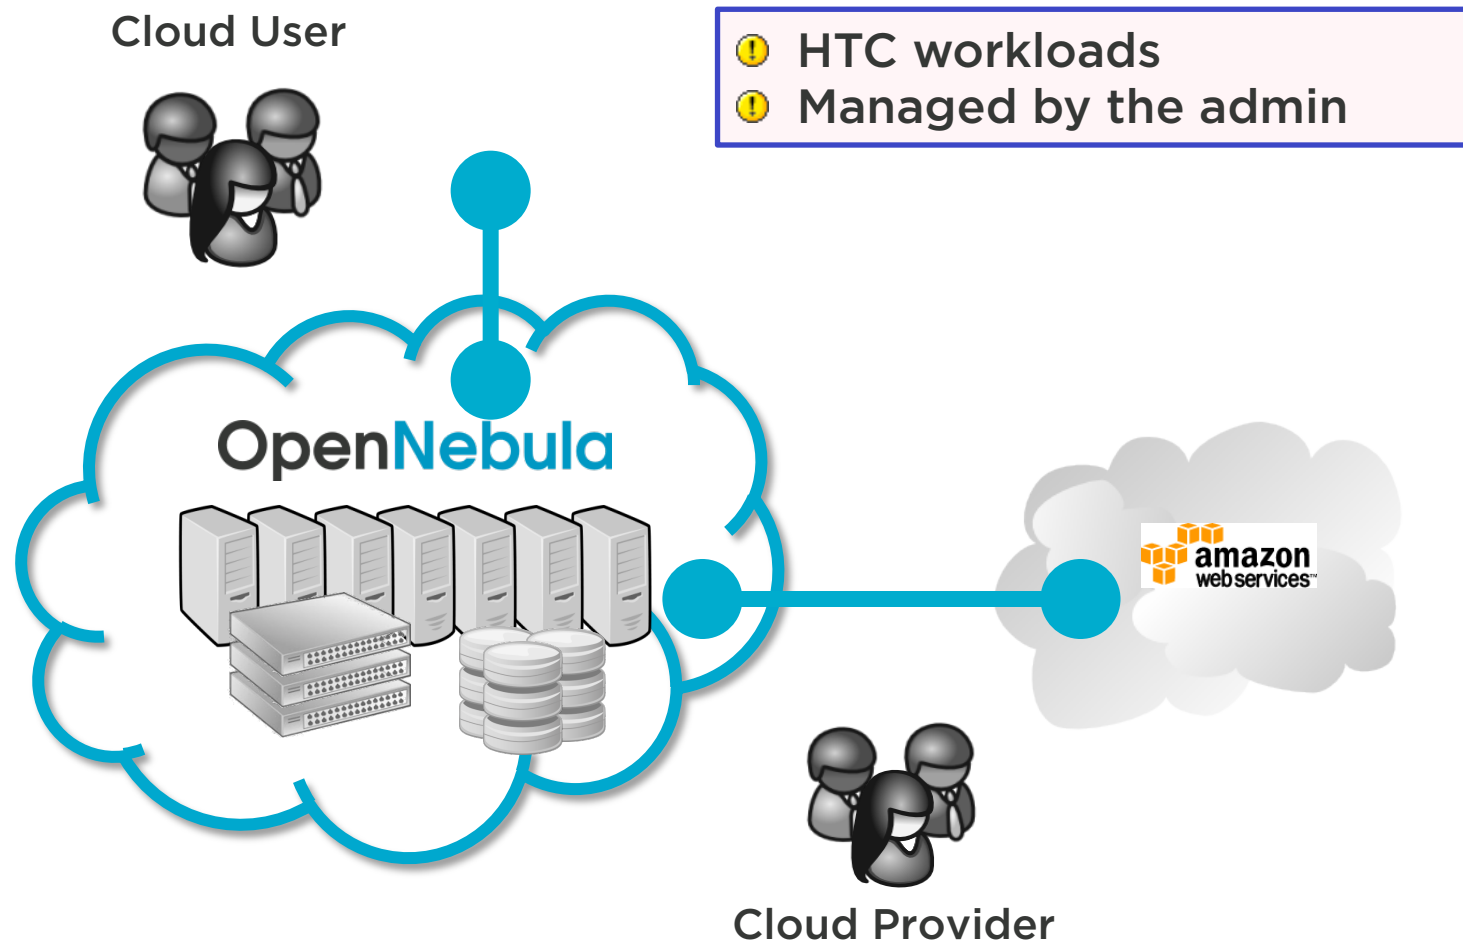
- ❗ Provide VOs or Departments with their own isolated cloud compartment
- ❗ Associated to clusters with specific Hw and Sw profiles, or security levels



## Hybrid Cloud Computing

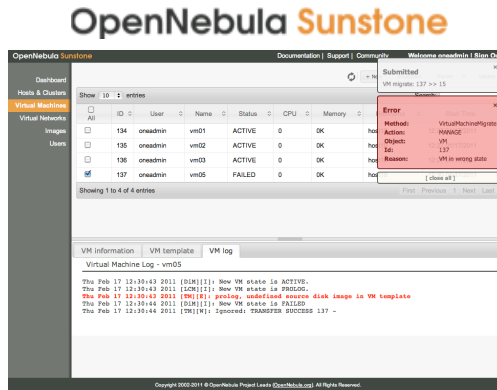
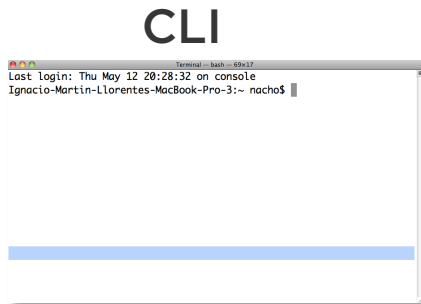
### Hybrid

- Extension of the local private infrastructure with resources from remote clouds
- **Cloudbursting** to meet peak or fluctuating demands



## How Can I Operate my Cloud Infrastructure?

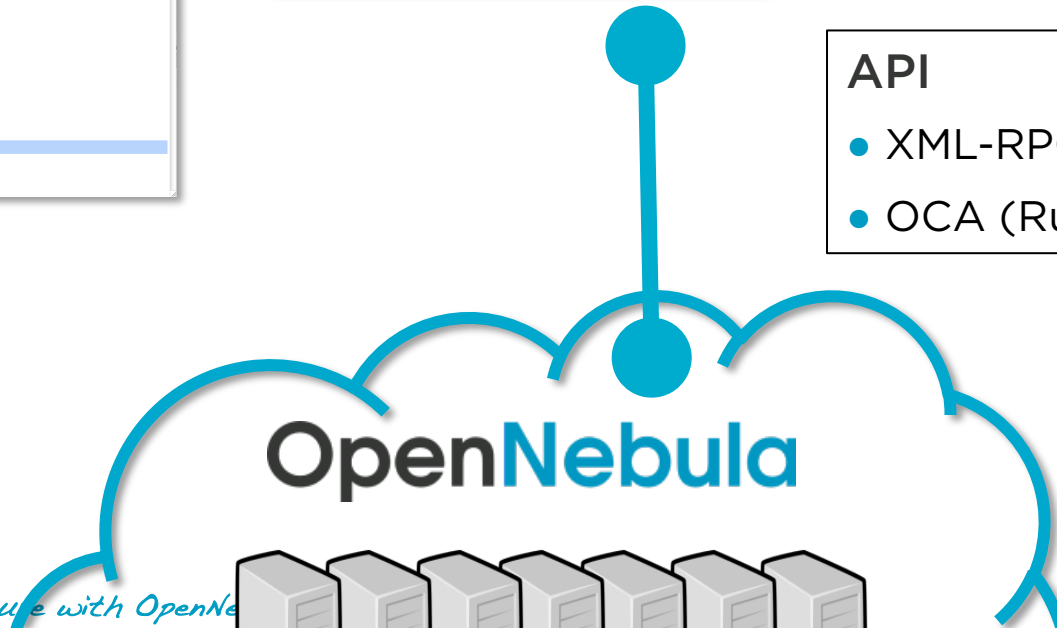
... programming APIs (create new tools and integrate), web interfaces (simplify operation), and command lined interface (create scripts)...



**Automatic/Elastic Management of Computing Service**

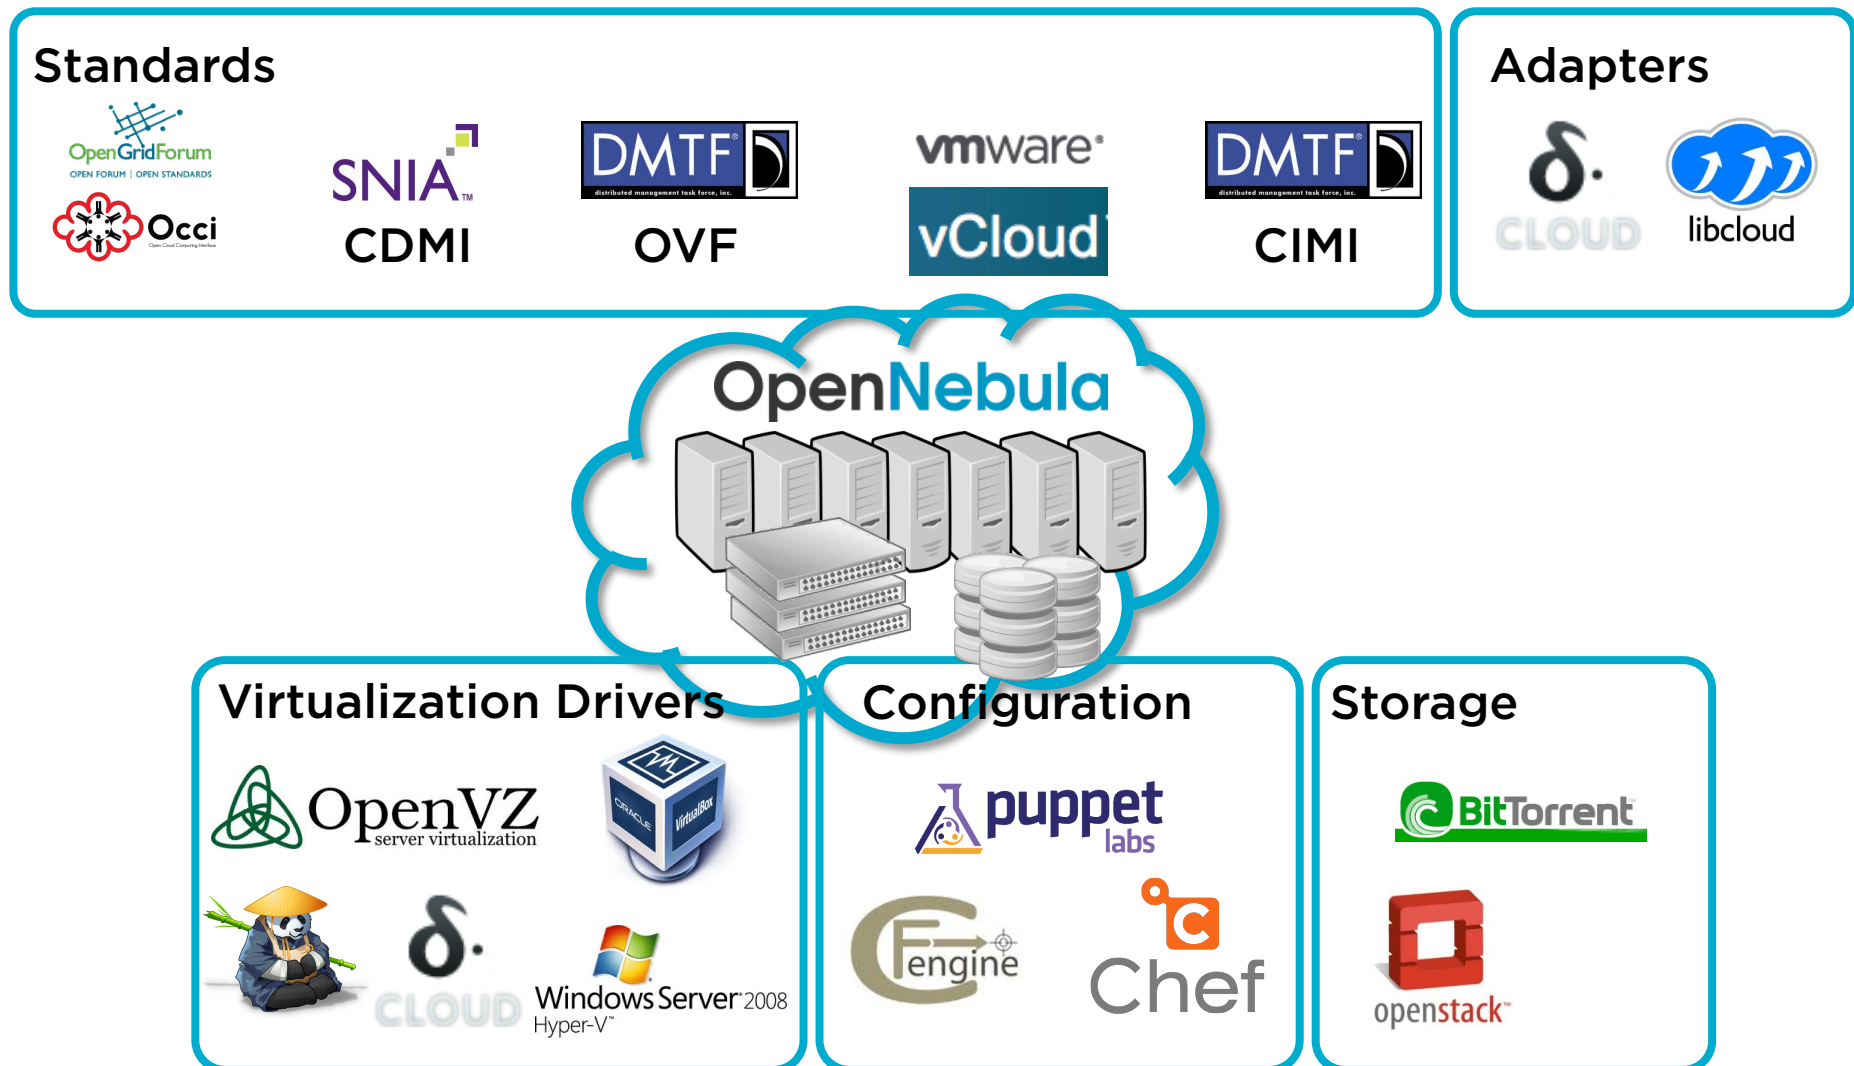
**API**

- XML-RPC
- OCA (Ruby, Java, Python)



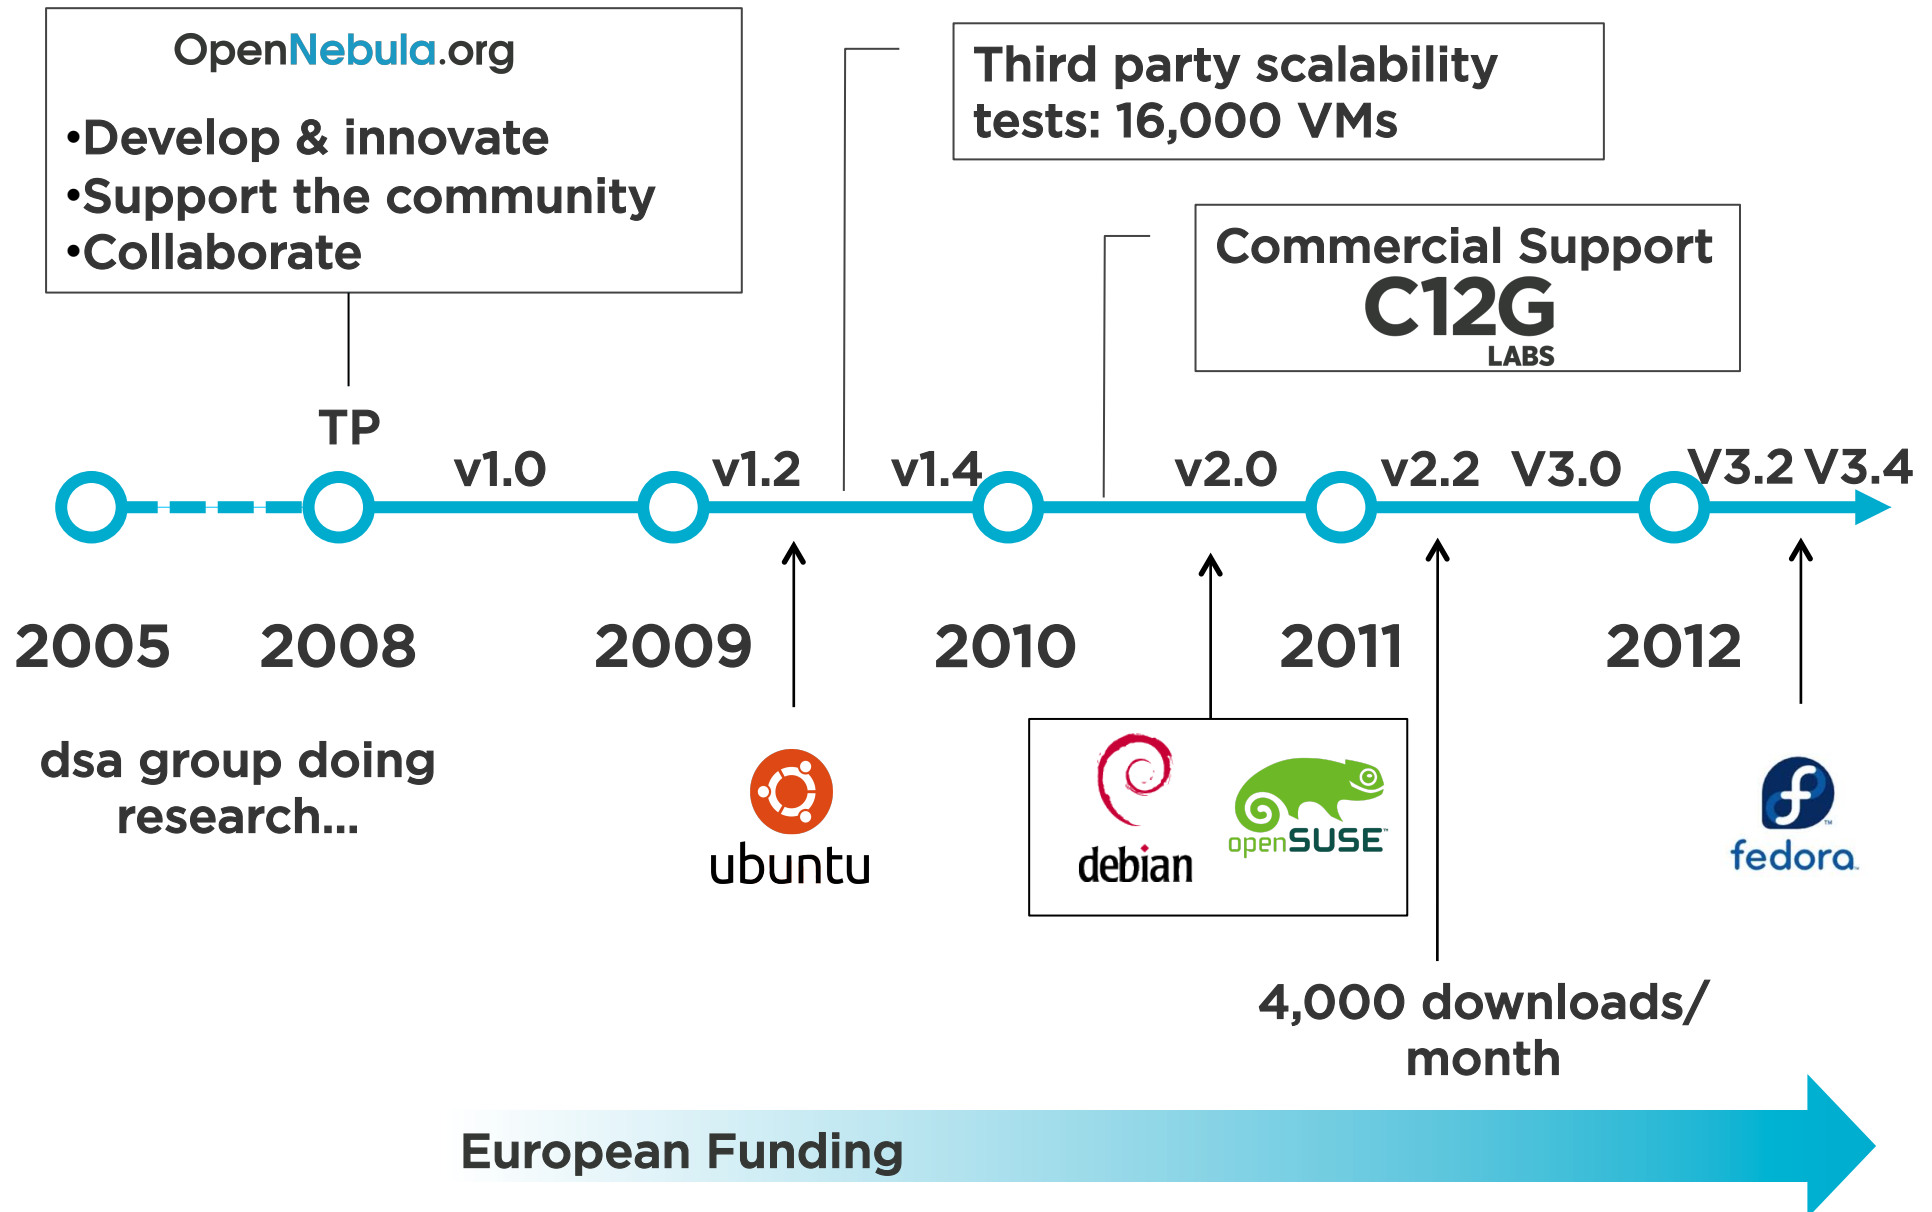
## The Ecosystem

... **truly open** (fully open-source, Apache license) and **adaptable** (modular and extensible)... because no two data centers are the same




# About the OpenNebula Project

*A Project Aimed at Building the Industry Standard Open Cloud Management Tool*



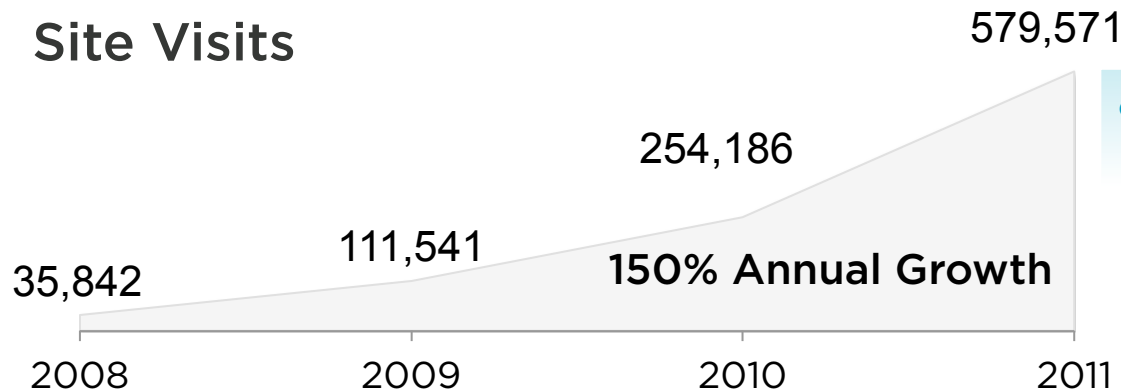
*A Quickly Growing Community More than Doubling Each Year*

## Downloads




- 900 downloads in the last week
- Linux distro and code repo

## Site Visits



- 15,300 visits and 194,000 page views in the last week

## Mails



- 800 registered users at present

Date: November 7<sup>th</sup>, 2011

*Open Cloud Enabler in the ICT Industry*

## Enabling Hosting Companies and Telcos to Offer Cloud Services



## Enabling Technology Companies to Offer Cloud Products



## Enabling Service Companies to Offer Cloud Consulting and Integration





*Open Cloud Enabler for Building and Operating Virtualized Data Centers*

## Industry



## Supercomputing Centers



## Research Centers






*Open Cloud Enabler for Building and Research and Innovation*

## Distributed Computing Infrastructures



## Research Projects



*We Will Be Happy to Answer Any Question*

The screenshot shows the OpenNebula.org website homepage. At the top right is the logo "OpenNebula.org" with the tagline "The Open Source Solution for Data Center Virtualization". Below the logo is a navigation bar with links: Home, About, Documentation, Software, Support, Community, Cloud, Dev, Blog, Wiki. The main content area features a large heading "Leading Cloud Management Solution" with the subtext "and it's free, fully open-source". Below this is a section titled "More than 4,000 Downloads per Month" and "Most Powerful Platform for Enterprise Data Center Innovation". A paragraph describes OpenNebula as the "open-source industry standard for data center virtualization". To the right of the main text is a sidebar with a list of links: "About OpenNebula.org", "What is OpenNebula?", "Why OpenNebula?", "Who is Behind?", "What are its Features?", "How Can I Try it?", and "How Can I Download it?". Below the sidebar is a section titled "Who is Using OpenNebula?". The bottom section features a grid of six benefits: "Open" (Fully open-source), "Adaptable" (Customizable to fit into your own data center), "Proven" (Many massive scale production deployments), "Powerful" (Advanced enterprise class functionality), "Interoperable" (Most popular cloud APIs and standard-based), and "No Lock-in" (Platform independent on major hypervisors). At the bottom, it says "Install OpenNebula from the Official Repo of your favourite Linux Distro." and shows logos for Debian, openSUSE, Ubuntu, and a "coming soon" logo. On the right side of the bottom section, there is a video player for "OpenNebula 3" with a list of features: "Advanced multi-tenancy", "Virtual data centers", and "Multiple zones", followed by "... and many more things" and a "Download now" link.