

Workshop GARR  
*Calcolo e Storage Distribuito*  
November 29-30, 2012 - Rome



# Cloud Computing in ENEA-GRID: Virtual Machines, Roaming Profile, and Online Storage

**Ing. Giovanni Ponti, Ph.D.**

*ENEA – UTICT-HPC*

[giovanni.ponti@enea.it](mailto:giovanni.ponti@enea.it)

***joint work with***

A. Rocchi, A. Colavincenzo, G. Giannini, A. Secco,  
G. Bracco, S. Migliori

- Introduction
  - Cloud Computing
  - ENEA-GRID
- Cloud Computing in ENEA-GRID
  - Previous Solutions
- New Proposal
  - Motivations
  - User Roaming Profile
  - Online Storage
  - Accessing the VMs
- Assessing the Cloud Infrastructure
  - Stress test
- Conclusion

**Pool of technologies and instruments able to export services in the net**

### ***Typical Services:***

- Storage/Archiving
- Data elaboration
- Remote Software
- Virtualization

### ***Advantages for users:***

- Ubiquitous data and resources access
- Device independance
- User-friendly interface and simple usage
- Saving of HW and SW cost and maintenance operations



## ENEA-GRID

ENEA computational centers connected each other by the GARR network

- More than 50Tflops (integrated power) and over 6000 cores (CRESCO HPC clusters)
- Integration of computational resources, storage systems, and monitoring tools
- A unified environment and homogeneous access for ENEA researchers
- Optimized resource usage
- Multi-platform systems for serial and parallel computations (HPC)



### ***Early Approach***

- Software and services in the GRID
- ENEA Virtual Labs  
(web access to software, tools, and docs related to a specific domain)
- Storage in GRID (OpenAFS)
- GPFS multisite  
(multicluster over WAN – *recently enabled*)

### ***Recent Approach***

- Testing ENEA-GRID infrastructure in supporting Virtual Machines customized on user requirements, employing OpenNebula as cloud platform.  
(GARR Conference - November 2011)
- OpenNebula is a versatile tool and supports interoperation of heterogeneous cloud and services (e.g., OCCl)

## Motivations to enhance Cloud services in ENEA-GRID

Providing **user-oriented VMs**, which can be used to support ENEA-GRID users in their common activities (not only regarding HPC)

### Requirements

- VMs should be *customized on user needs*
- Providing support to **legacy applications**
- VMs should **access user data** in ENEA-GRID
- There should be an **easy way to share and transfer data** from and to the VMs
- Store **user personalizations and environment** (profile)

Infrastructure **also** suitable to provide **service-oriented VMs**

## Newest Approach

To enhance virtual machine supply by introducing ***user roaming profiles*** and ***online storage*** for data exchange

### ***Advantages***

- Handling user roaming profiles (both for Windows and Linux systems) allows to save user personalizations and settings on his VMs
- Online storage facility allows to easily exchange data from and to VM
- Both roaming profile and online storage exploit AFS as storage system in cloud environment and Kerberos5 for authentication

# User Roaming Profiles for VMs



- VMs join ENEA.IT domain
  - Authentication with Krb5
- AFS is employed to store user profile data
  - One profile for each VM template available
- Users login into VMs with their ENEA-GRID credential
  - After login, users can access to their AFS home in ENEA-GRID to share data among VMs
  - *No need for creating local accounts in VMs!!!*
- VMs can be destroyed after usage without losing user data, as they are stored on AFS
  - Optimized usage of infrastructure resources

# Online storage – OKBox



- A typical challenge in cloud systems and in VM context concerns the ***data exchange issue***
  - Users want to transfer data from and to their VMs in a easily way
- ENEA provides an own solution by proposing **OKBox**

## OKBox

Service for distributed storage system which allows to share data easily in ENEA-GRID cloud



### Main features:

- It exploits IaaS strategy and “Always and Anywhere On” cloud paradigm
- Web access (pervasive) to distribute resources in AFS
- Drag-and-drop file upload
- Possibility of sharing data among ENEA-GRID users and also among other people (non ENEA-GRID users)
- Security policies and privacy preserving mechanisms (no use of third-part storage solutions)

# OKBox – Login Interface



## ENEA OKBOX Login

The screenshot shows the ENEA OKBOX Login interface. A red box highlights the 'Login' and 'Password' input fields. A red circle with a dot is placed over the 'Login' field, with an arrow pointing to the text 'Login credentials'. Another red circle with a dot is placed over the 'Password' field, with an arrow pointing to the text 'ENEA-GRID user flag'. Below the input fields, the text 'I'm an ENEA-Grid user' is followed by a checkbox, which is also highlighted with a red box and a red circle with a dot, with an arrow pointing to the text 'Access points'. Below this, the text 'Take me to:' is followed by three radio button options: 'My Own Box', 'A Project's Box', and 'Another User's Box'. A red oval is drawn around these three options, with an arrow pointing to the text 'Own box', 'Project box', and 'Another user box'. At the bottom left, there is a 'Log-in' button.

• Login credentials

• ENEA-GRID user flag

• Access points

– Own box

– Project box

– Another user box

# OKBox – User Interface



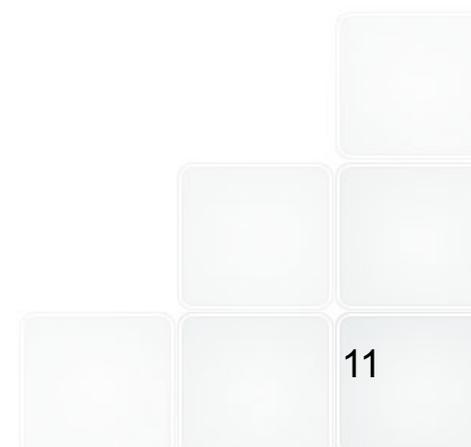
Logged in as **gponti** Logout

Up one level   New folder   Share the current folder  
Delete   Rename   Generate link   Upload

/afs/enea.it/por/user/gponti/public\_okbox/

Filename	Type	Last Modified	Size
<input type="checkbox"/> casa	directory	Nov 28 00:20	-
<input type="checkbox"/> prova	directory	Nov 23 10:56	-
<input checked="" type="checkbox"/> 2012_11_30_PONTI_ENEA_GRID_CLOUD.pdf	file	Nov 28 00:46	1026662
<input type="checkbox"/> 2012_11_30_PONTI_ENEA_GRID_CLOUD.ppt	file	Nov 28 00:47	1268736

- Data area (files and dir)
- User operations



# Accessing the VMs



## Ways to access to the VMs

- ENEA-GRID FARO Web Portal (ad-hoc interface)

The screenshot shows the ENEA-GRID FARO Web Portal interface. At the top, there is a navigation bar with icons for Refresh, Connect (VNC), Connect (RDP), Stop VM, Start VM, Delete VM, and Sunstone Portal. Below this is a table listing VMs with columns for ID, Status, Name, Protocol type, Hosted on, and Address. The table contains the following data:

ID	Status	Name	Protocol type	Hosted on	Address
75	✓ (runn)	TrixB0x	vnc	crescolx022::23990	192.107.70.83
168	✓ (runn)	LSF Test	vnc	crescolx020::6068	192.168.13.5
169	✓ (runn)	LSF Test	vnc	crescolx022::6069	192.168.13.6
178	✓ (runn)	test	vnc	crescolx020::6078	192.168.13.7
185	✓ (runn)	win_ale_32_test	vnc	crescolx021::6085	192.168.13.230

Below the table is a 'VM Creation' section with a 'Template List' containing the following items: WinTest64, Ubuntu 10.04.4 PublicNet, Ubuntu 10.04.4 Vlan1, WinXp\_64, TrixB0x, Windows\_7\_64, WinXp\_32, WinXP32\_OLDNebula, Centos\_5.3, Ubuntu 10.04.4 Vlan1 TestDep, and Win\_XP\_32\_PONTI. There are buttons for 'Refresh the Template List' and 'Create another VM with the selected template'.

- Simple interface wrapping OpenNebula APIs to execute most common operations on VMs (e.g., create, stop, delete)
- Accessing protocols VNC and RDP
- OpenNebula ACLs on users and group employed for better managing templates and physical resources

- OpenNebula Sunstone web portal (for advanced users)

## ***Cloud Infrastructure:***

- 3 IBM x3850/x3950-M2 nodes with:
  - 4 Xeon Quad-Core Tigerton E7330 (2.4GHz/1066MHz/6MB L2)
  - 32 Gb RAM
- Total amount of physical resources:
  - 48 cores
  - 96 Gb RAM
- 1 node running OpenNebula processes, scheduler, and Sunstone Web Portal

## ***VM Template Info:***

- *Operating System:* Windows XP 32bit SP3
- *CPU:* 1 physical CPU per VM
- *RAM:* 1 Gb per VM

## Purpose

Assessing the responsiveness of the VMs

### ***Two-stage test:***

**1. TEST 1:** Standard loaded infrastructure  
(physical resources yet available)

- Up to 48 VMs

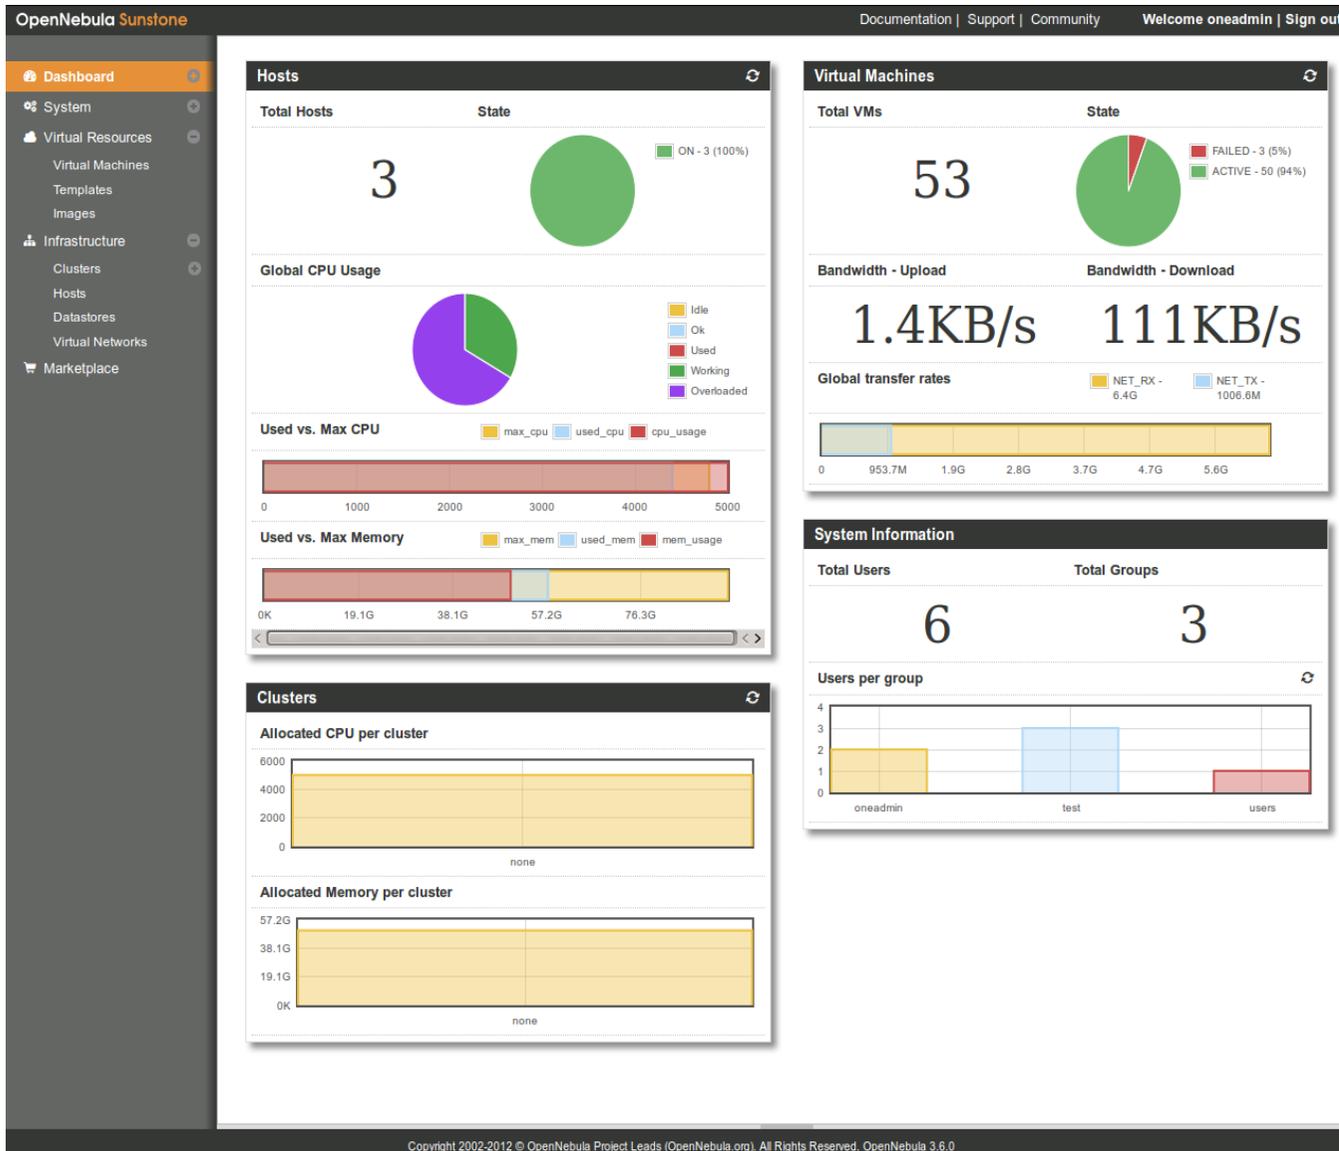
**2. TEST 2:** Overloaded infrastructure  
(physical resources partially overloaded)

- More than 48 VMs

### ***VM Status:***

- Each VM has its CPU fully busy (100%)
- Factorization of large integers (process task executed)

# Testing the Cloud Infrastructure (3)



## Results:

- **TEST 1:** System highly responsive
- **TEST 2:** System yet responsive and VMs usable for common user activities (e.g., web browsing, document editing, etc...)

- We presented the cloud infrastructure in ENEA-GRID based on OpenNebula
- We aimed to provide disposable VMs for user purposes

## ***Key features of our proposal:***

- User exploited VM integration in ENEA-GRID by means of:
  - Centralized roaming profile;
  - Online data exchange storage solution (ENEA OKBox)
- Easy access via FARO Web Portal
- Infrastructure supporting our objective (stress test)

**Ready to be released for ENEA-GRID users**



# Thanks!

**Ing. Giovanni Ponti, Ph.D.**

*ENEA – UTICT-HPC*

[giovanni.ponti@enea.it](mailto:giovanni.ponti@enea.it)



***People in ENEA-GRID/CRESCO***



D. Abate, F. Ambrosino, G. Aprea, T. Bastianelli, F. Beone, M. Caporicci, M. Chinnici, A. Cucurullo, P. D'Angelo, A. Della Casa, M. De Rosa, A. Funel, G. Furini, D. Giammattei, S. Giusepponi, R. Guadagni, G. Guarnieri, A. Italiano, A. Mariano, A. Perozziello, A. Petricca, G. Mencuccini, C. Mercuri, P. Ornelli, S. Pecoraro, S. Pierattini, S. Podda, F. Poggi, A. Quintiliani C. Sciò, F. Simoni

***ENEA-GRID/CRESCO***

<http://www.cresco.enea.it/>