



Status of Cloud Computing Project @ IFJ PAN

Henryk Palka
Institute of Nuclear Physics PAS, Krakow



The Project

- Funded by European Commission and Polish Ministry of Science and Education (Innovative Economy, National Cohesion Strategy) for 2009-2013, budget ~10M PLN
- Launched in October 2009 (written up in 1Q 2009, submitted in April 09, accepted in Sept 09)
- Acronym: CC1 (CracowCloud1)
- The project goals:
 - 'to diminish computational barriers in science and economy by employing elastic computing technology'
 - i.e.: easy access, run multi-operating-system applications



The Project Goals

- integrate computing infrastructure in-house (private cloud)
- integrate infrastructures of research institutions (distributed public cloud):
 - Krakow's Technical University
 - University of Kielce (100km north)
 - University of Rzeszow (180km east)
 - + CYFRONET (within PL GRID+)
- increase computing efficiency, lower operational costs
- attract IT startups with good ideas (and no capital usually), show the way to use commercial clouds



The Team

➤ Developers:

Krzysztof Danielowski
Rafal Grzymkowski
Maciej Kruk
Maciej Nabozny
Tomasz Sosnicki
Przemyslaw Syktus
Piotr Wojcik
Milosz Zdybal

➤ Hardware infra. support:

Zbigniew Natkaniec
Bartlomiej Zabinski

➤ Software support:

Janusz Chwastowski
Andrzej Olszewski

➤ Project Coordinator: Mariusz Witek

➤ Project Leader: HP



Software

➤ Timeline of the system development:

Phase1 (→ 1Q 2011): - evaluation of OpenSource Cloud Managers
- evaluation of virtualization engines,
- development of tools and interfaces,
- establish functional system for local users

Phase2 (→ 1Q 2012): distributed Cloud

Phase3 (→ 1Q 2013): - production system (also for Belle II use)
- admitt IT startups

➤ Phase1 till now: - 3 testbeds, 2 blades each (16 cores)
- evaluation teams (1/testbed):
Eucalyptus, OpenNebula, Nimbus



Evaluation Result (Eucalyptus, ONe, Nimbus)

- **Eucalyptus**: powerfull, many features (interf. to EC2, multi-clusters), put aside because:
 - non-modular (difficult to modify, expand)
 - lack of control on a user
 - complex and large logs (difficult to turn off)
 - OpenSource version more and more behind the comercial version
- **Nimbus**: the project closest to HEP, uses GRID elements (authentication, GRIDftp...GLOBUS) discarded because:
 - poorly documented, difficult to install, couldn't get working many of the claimed fetures
 - requires running the own CA
- **OpenNebula**: modular, light, has hooks , many missing functionalities **our choice** because: it looks like it is a good software, easy to expand



Software

➤ Status as of today:

- CloudManager selected: OpenNebula
- XEN seems to be our preferred virtualizer (KVM still investigated)
- Our own extensions and improvements to ONe have been develop^{ed}, (next talks)
- Went even beyond that: it could be that we will abandon ONe ! (providing only a plug-in to ONe), since our own CM looks promising, smarter than ONE
→ today's talks by Tomasz, Maciek, Przemek



Hardware: Server Room

➤ Tender requirements:

- construction works: area → 60m²
- 2 independent power lines 250kW
- technical floor (rised)
- reduntant cooling : 2 cold air units, 2x80kW
- reduntant UPS: 4x40kVA
- fire extinguisher system: gas, INERGEN

➤ Tender announcement: Nov 2009

➤ Tender concluded: Mar 2010

➤ Cost: ~1M PLN

➤ Delivery: 70 days

➤ Status: nearly ready



IFJ PAN Server Room: Sept 2009



Server Room: April, 2010



Server Room: May, 2010



Hardware: Blades, Storage, Networking

➤ Tender requirements:

- Blades: ≥ 4 cores/CPU, ≥ 500 cores, 2GB/core, LV (with specs pointing to Intel Westmere 6 cores), 140GB HD SAS, 2xEth, 2xIB QDR 4X, Eth & IB switches
Racks upgradable to ice-water cooling
- Storage: 45TB, SSD cache (SUNFire)
- Networking: IB, 10GB switch, dozen of 1Gb switches, WiFi AP, Firewalls

➤ Tender announcement: Mar 2010

➤ Tender concluded: June 2010 (with many troubles)

➤ Cost: ~2.5M PLN

➤ Delivery: end of Aug 2010

➤ Status: nearly ready

➤ Next tender, purchase of the second batch: 2011





Grid vs Cloud

Grid



Cooperative harvesting

Cloud



Contract harvesting



Cloud Computing

➤ Nice term : grips the imagination (→marketing),
but doesn't explain anything

My preference : **elastic computing**
(however do not pretend to tame
spread of cloud)

'elasticity':

- Run multi-opsys applications
- Configure hardware in min^{ts}
- Workload overflow to other clusters in the cloud

.....

➤ Central concept: **VIRTUALISATION**

