



the globus alliance

www.globus.org

Nimbus  
or  
an Open Source Cloud Platform  
or  
the Best Open Source EC2 No  
Money Can Buy ;-)

Kate Keahey

Tim Freeman

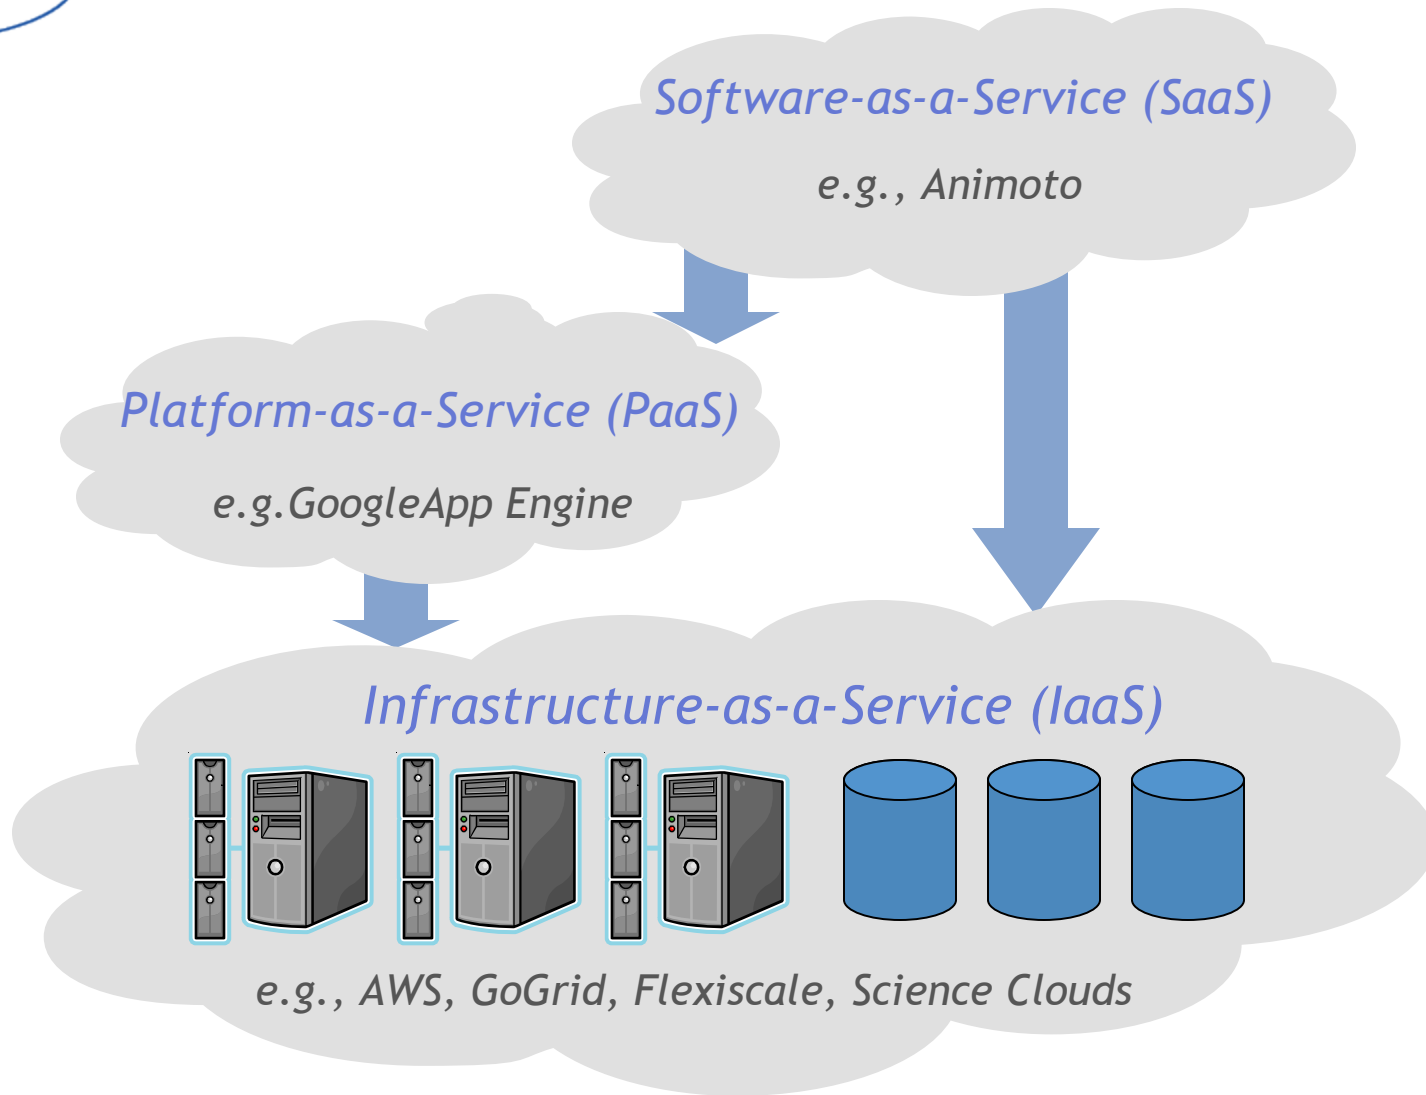
*(keahey@mcs.anl.gov)*

University of Chicago

Argonne National Laboratory



# What Are Clouds?





# What is Nimbus?

- An extensible open source Infrastructure-as-a-Service implementation
  - ◆ Turns your cluster into a cloud
- Why open source IaaS?
  - ◆ **Experiment and use:** make your own cloud or configure a private cloud
  - ◆ **Customize:** try new things, make the IaaS paradigm work for your application domain
- Our particular interest in customization: scientific computing




# Nimbus Features

- Cloud computing infrastructure
  - ◆ WSRF and EC2 interfaces
  - ◆ Xen implementation (KVM in preparation for release)
  - ◆ Launches flexibly defined groups of VMs and configures their networking
- Can be configured to use familiar schedulers like PBS and SGE to manage VMs
  - ◆ The workspace pilot
- Launches self-configuring virtual clusters
  - ◆ The context broker
- Defines an extensible architecture
  - ◆ And has been extended by various projects

**News Flash: Ian Gable from UVIC becomes a Nimbus committer!**



# 5-click Guide to Nimbus





# Science Clouds

**News Flash: the Masaryk cloud officially available online today!**

## Introduction to Nimbus @ UC

The Ur  
in the 1  
Univer:

**ACIS** Advanced Computing and Information Systems Laboratory



Home | People | Research | Workspace | Education |

**PURDUE**  
UNIVERSITY



**TeraGrid**

### Cloud Computing Resource:

Purdue provides a cloud computing testbed "Wispy" to TeraGrid users. It contains 100 half gigabytes of available memory. 100 VCPUs



**Kupa -- Cloud installation in MetaCenter**




# Science Clouds: Goals

- Make it easy for scientific projects to experiment with cloud computing
  - ◆ Can cloud computing be used for science?
- Customize software in response to the needs of scientific projects
  - ◆ Start with EC2-like functionality and evolve to serve scientific projects: virtual clusters, diverse resource leases
  - ◆ Federating clouds: moving between cloud resources in academic and commercial space



# Who Runs on Science Clouds?




*Project diversity: Science, CS, education, build&test...*





# CloudBLAST



- CS research: investigate latency-sensitive apps, e.g. hadoop
- Need access to distributed resources, and high level of privilege to run a ViNE router
- Virtual workspace: ViNE router + application VMs
- *Paper: "CloudBLAST: Combining MapReduce and Virtualization on Distributed Resources for Bioinformatics Applications" by Andréa Matsunaga, Maurício Tsugawa and José Fortes, accepted to eScience 2008.*



- STAR: a high-energy physics experiment
- Needs resources **with the right configuration**
  - ◆ Complex environments: correct versions of operating systems, libraries, tools, etc all have to be installed.
  - ◆ Consistent environments: require validation
- Solution: deploy a virtual OSG STAR cluster
  - ◆ OSG CE (headnode) plus STAR worker nodes: SL4 + STAR conf
- Requirements
  - ◆ One-click virtual cluster deployment
  - ◆ Migration: Science Clouds -> EC2
    - Last September: EC2 runs of up to 100 nodes (production scale, non-critical codes)
    - Testing for critical production deployment
- Work by Jerome Lauret, Doug Olson, Leve Hajdu, Lidia Didenko



# Infrastructure Testing

- Project: tests of various Globus components on different platforms for correctness and scalability
- Need short-term but flexible access to diverse platforms
- Workspaces: Globus 101 + others
- Work by various members of the Globus Toolkit (Tom Howe, John Bresnahan, Stu Martin, Martin Feller)
- Resulted in provisioning a private Nimbus cloud for Globus



# Alice HEP Experiment at CERN



## MonALISA Repository for ALICE



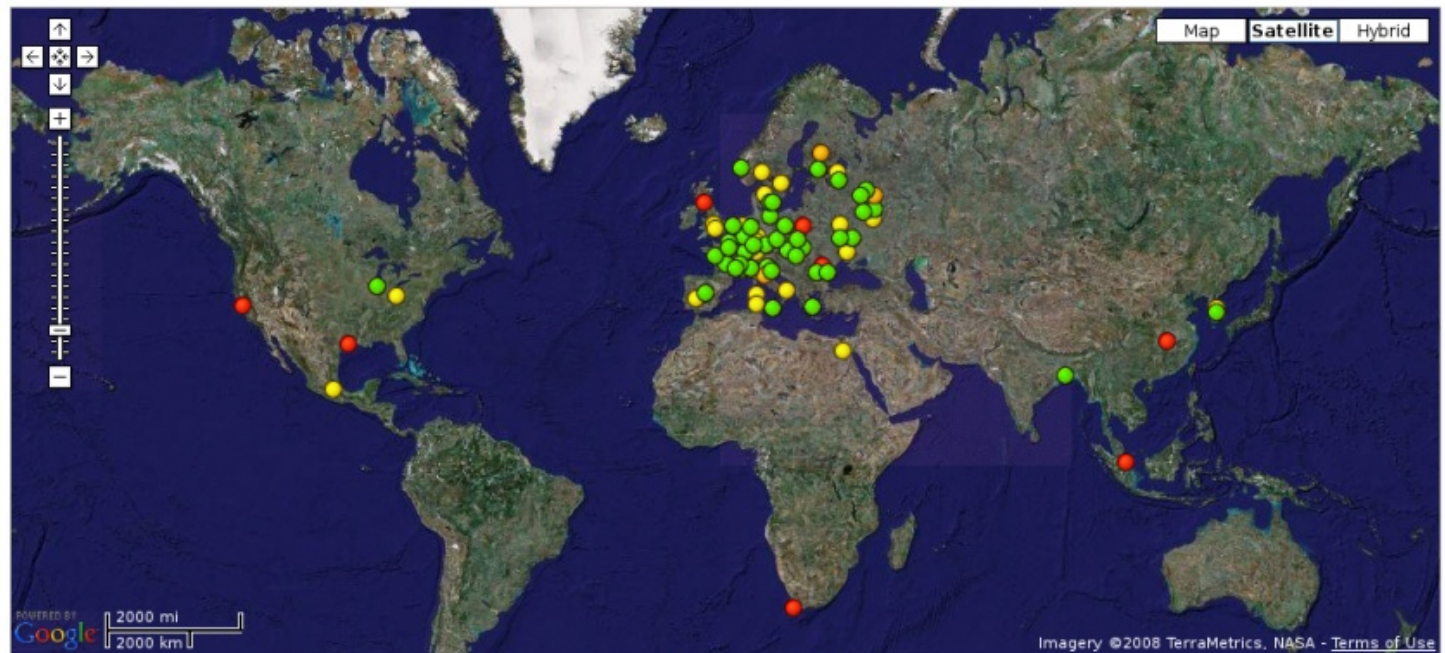
Repository Home Administration Section ALICE Reports Events XML Feed Firefox Toolbar MonALISA GUI

ALICE Repository

- ALICE Repository
- Google Map
- Shifter's dashboard
- Running trend
- Production info
- Job Information
- SE Information
- Services
- Network Traffic
- FTD Transfers
- CAF Monitoring
- SHUTTLE
- LCG exp. monitoring
- Build system
- Dynamic charts

close all

This page: bookmark, URL



● Running Jobs ● ML Service Down ● No Active Jobs ● ML Service Down & no running jobs

Map options

Show xrootd transfers  Show site relations

Jump to: Europe North America South America Asia World Save position and options

## Collaboration with the CERNVM project



# Other Projects

- Evaluating a cloud from user's perspective
  - ◆ *Paper: "Exploration of the Applicability of Cloud Computing to Large-Scale Scientific Workflows", C. Hoffa, T. Freeman, G. Mehta, E. Deelman, K. Keahey, SWBES08: Challenging Issues in Workflow Applications*
- Economics
  - ◆ *Paper: "Testing Different Imitation Strategies in PD Game on Networks", Svarc, P. and N. Svarcova, Journal of Economic Interaction and Coordination (JEIC-D-08-00034).*



[Home](#)

[News](#)

[Documentation](#)

[Community Resources](#)

[Download](#)

[Publications](#)

[Science Clouds](#)

## Have a comment?

If you'd like to share your experiences with Nimbus, please sign our guestbook:

Thanks for your feedback!

Send

## Our Guestbook: see what users are saying!

---

*"Xen was the obvious choice for the virtualization part, but it took me a fair amount of time to find a piece of software to manage the deployment of the VMs. I came across Globus virtual workspaces (Nimbus) and the Globus Toolkit. This was by far the best solution for deploying and managing the developer environments."*

Scott Haskell, AdBrite.com, October 29, 2008

---

*"\*\*VERY COOL\*\* .. you have done really nice work."*

Nicholas Karonis, Northern Illinois University, October 20, 2008

---

*"I've been around HPC for 10 years and this is without question the coolest thing I've seen in a long time!"*

Ron Price, The University of Utah, October 16, 2008



[Home](#)[News](#)[Documentation](#)[Community Resources](#)[Download](#)[Publications](#)[Science Clouds](#)[Clouds  
Overview](#)[Marketplace](#)[Client  
quickstart](#)[One-click  
clusters](#)[Cluster guide](#)[Appendix](#)

## Science Clouds

---

### Available Science Clouds:

---

**To obtain access to a cloud, please contact the administrators of the clouds below.** Read cloud configuration notes carefully to ensure that you end up with the right software and configuration.

- [Nimbus @ University of Chicago](#)
- [Stratus @ University of Florida](#)
- [Wispy @ Purdue University](#)
- [Kupa @ Masaryk University](#)

To get started follow the instructions in the [cloud client quickstart guide](#). Please subscribe to [workspace-user@globus.org](mailto:workspace-user@globus.org) for updates on cloud status ([subscribe](#)).

- Available to scientific projects
- Send us mail to get access
- Do the quickstart: you should be running in 15 minutes!
- Move on to create virtual clusters and virtual Grids



# The Last Slide

- Nimbus is an extensible, easy-to-use, open source tool for configuring clouds
- What has our impact been?
  - ◆ Utilization, time used per project, etc.
  - ◆ Scientific results, papers written and in preparation, ongoing discussions
- Another kind of impact: we are doing things we could not do before
  - ◆ Deploying network routers on remote platforms
  - ◆ Easily finding the right environment in distributed environment
  - ◆ Provisioning resource when we need them
- We're learning what's possible