Cloud Computing with Nimbus

February 2009

Kate Keahey

(keahey@mcs.anl.gov)

University of Chicago

Argonne National Laboratory

Cloud Computing

elasticity
computing on demand

capital expense
⇩
operational expense

SaaS
Software-as-a-Service

PaaS
Platform-as-a-Service

IaaS
Infrastructure-as-a-Service

The Quest Begins

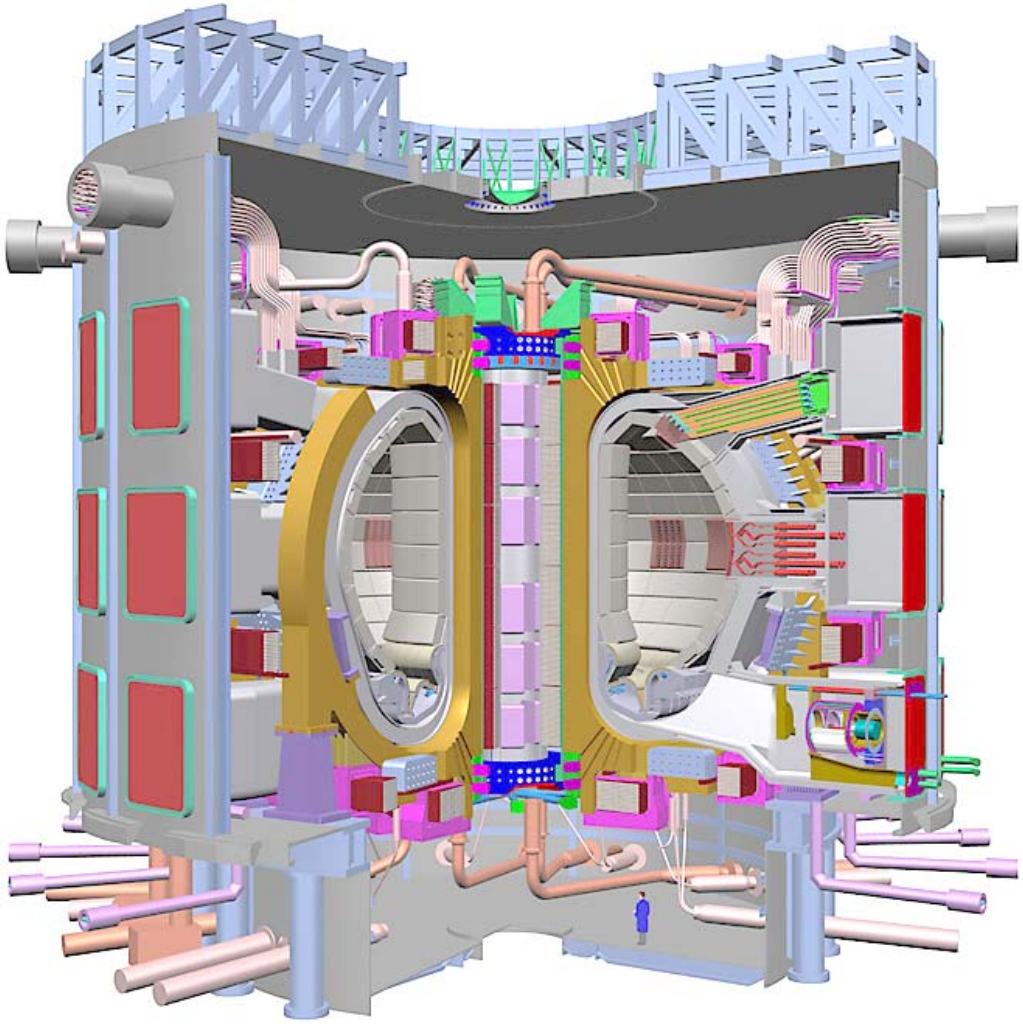

- Code complexity
- Resource control

“Workspaces”

- Dynamically provisioned environments
  - Environment control
  - Resource control
- Implementations
  - Via leasing hardware platforms: reimaging, configuration management, dynamic accounts…
  - Via virtualization: VM deployment

A Brief History of Nimbus

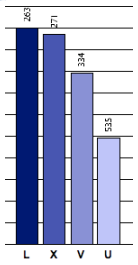

Xen released

EC2 goes online

First STAR production run on EC2

Nimbus Cloud comes online

2003

2006

2009

Research on agreement-based services

First WSRF Workspace Service release

EC2 gateway available

Support for EC2 interfaces

Context Broker release

Nimbus Overview

- Goal: cloud computing for science
  - Open source, extensible IaaS implementation
    - A platform for experimentation with features for scientific needs and interoperability
    - Set up private clouds (privacy, expense considerations)
    - Workspace Service
  - Orchestration tools
    - Focus on end-to-end picture
    - Context Broker, gateway
- http://workspace.globus.org/

The Workspace Service

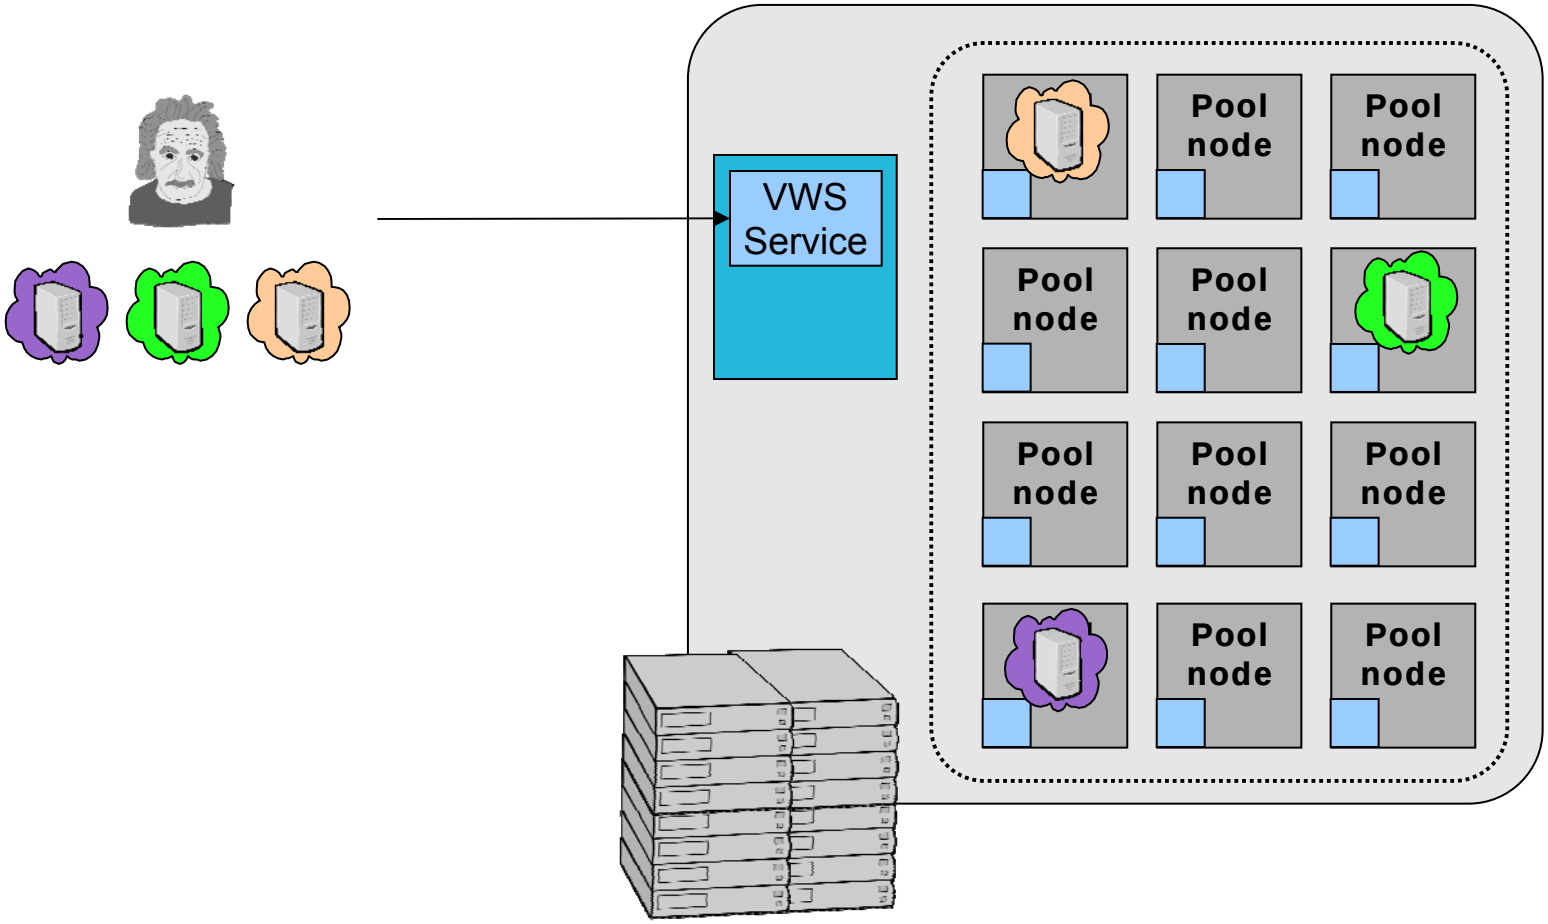

The Workspace Service

The workspace service publishes information on each workspace

Users can query this information to find out things about their workspace (e.g. what IP the workspace was bound to)

Users can interact directly with their workspaces the same way the would with a physical machine.

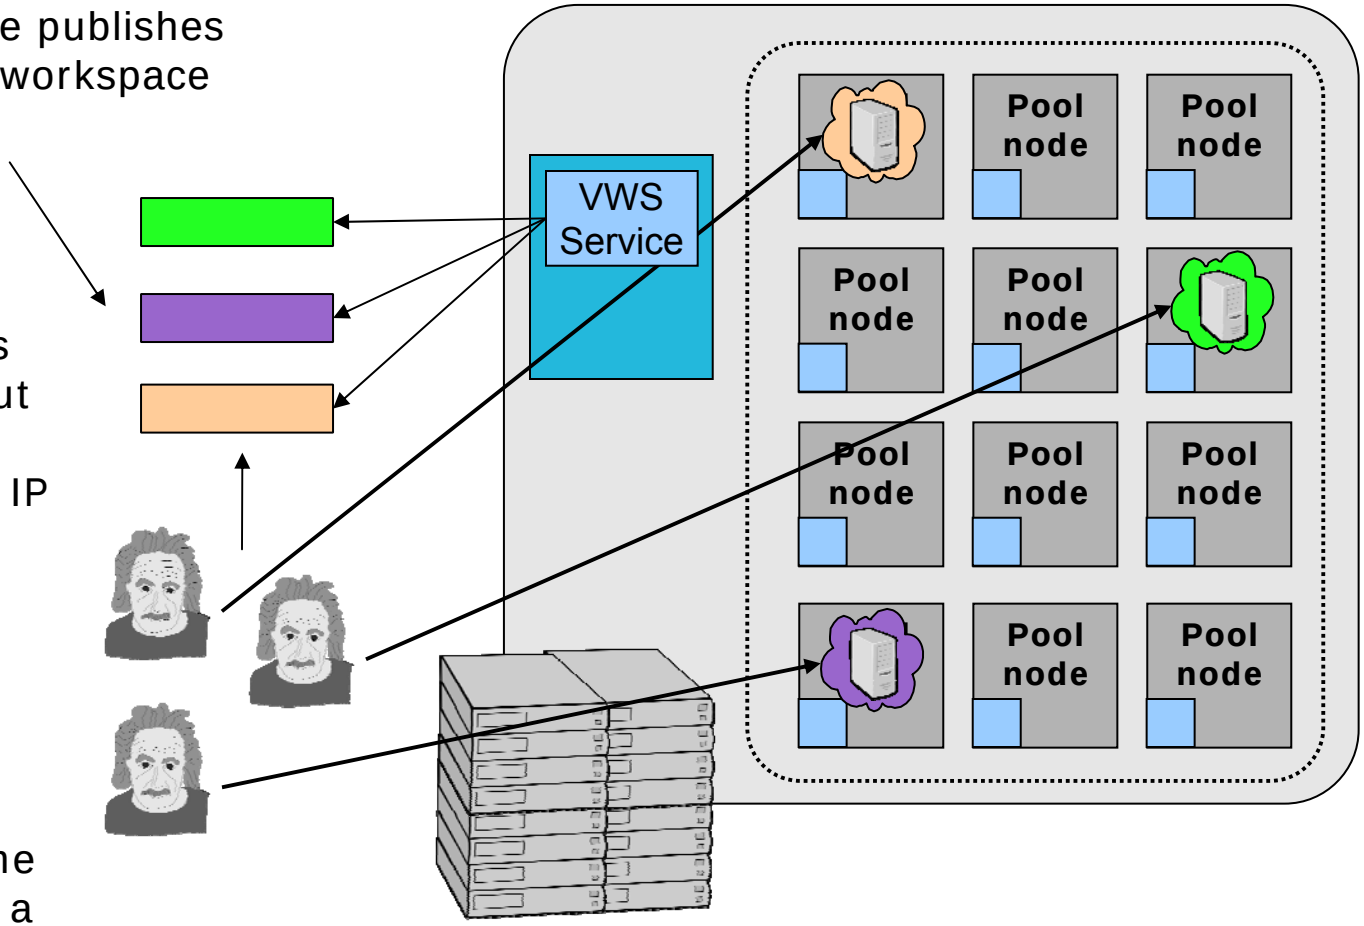

Trusted Computing Base (TCB)

Workspace Service Interfaces and Clients

- Web Services based
- Web Service Resource Framework (WSRF)
  - GT-based
- Elastic Computing Cloud (EC2)
  - Supported: ec2-describe-images, ec2-run-instances, ec2-describe-instances, ec2-terminate-instances, ec2-reboot-instances, ec2-add-keypair, ec2-delete-keypair
  - Unsupported: availability zones, security groups, elastic IP assignment, REST
- Used alongside WSRF interfaces
  - E.g., the University of Chicago cloud allows you to connect via the cloud client or via the EC2 client

Security

- GSI authentication and authorization
  - PKI credential required
  - Works with Grid proxies
  - VOMS, Shibboleth (via GridShib), custom PDPs
- Secure access to VMs
  - EC2 key generation or accessed from .ssh
- Validating images and image data
  - Collaboration with Vienna University of Technology
  - *Paper: Descher et al., Retaining Data Control in Infrastructure Clouds*

Networking

- Network configuration
  - External: public IPs or private IPs (via VPN)
  - Internal: private network via a local cluster network
- Each VM can specify multiple NICs mixing private and public networks (WSRF only)
  - E.g., cluster worker nodes on a private network, headnode on both public and private network

The Back Story

Workspace **WSRF front-end** that allows clients to deploy and manage virtual workspaces

Workspace back-end:

Resource manager for a pool of physical nodes Deploys and manages Workspaces on the nodes

Each node must have a VMM (Xen) installed, as well as the **workspace control** program that manages individual nodes

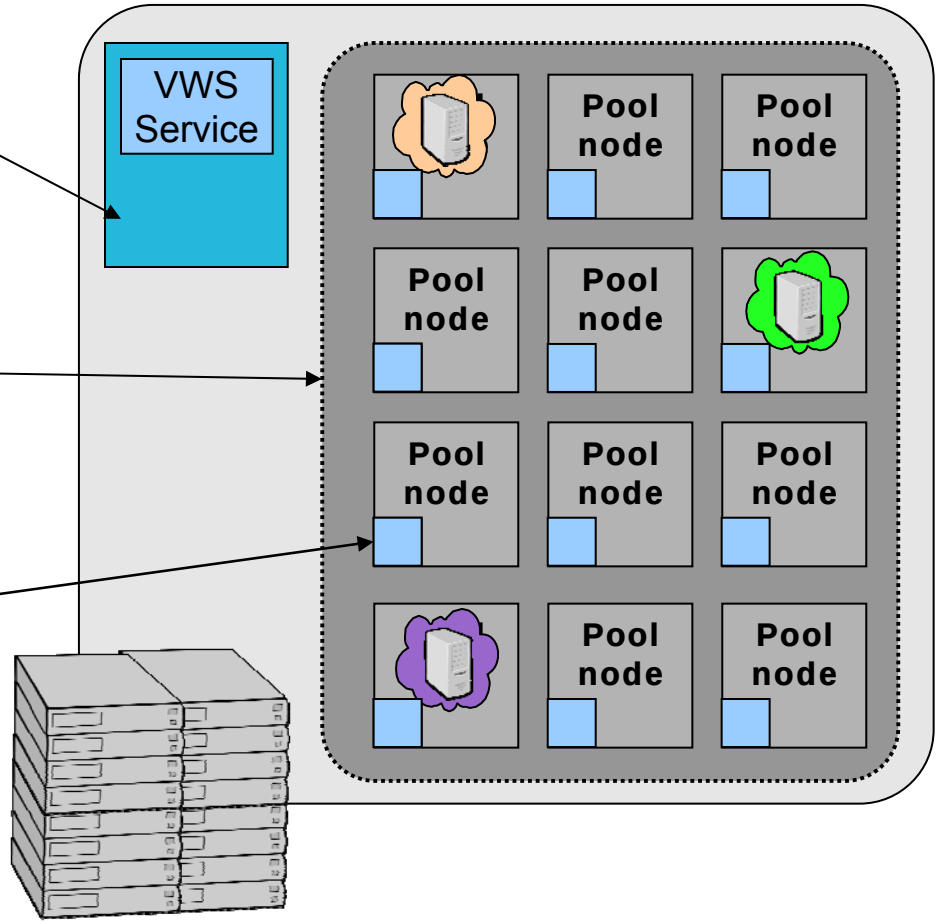

Trusted Computing Base (TCB)

Workspace Components

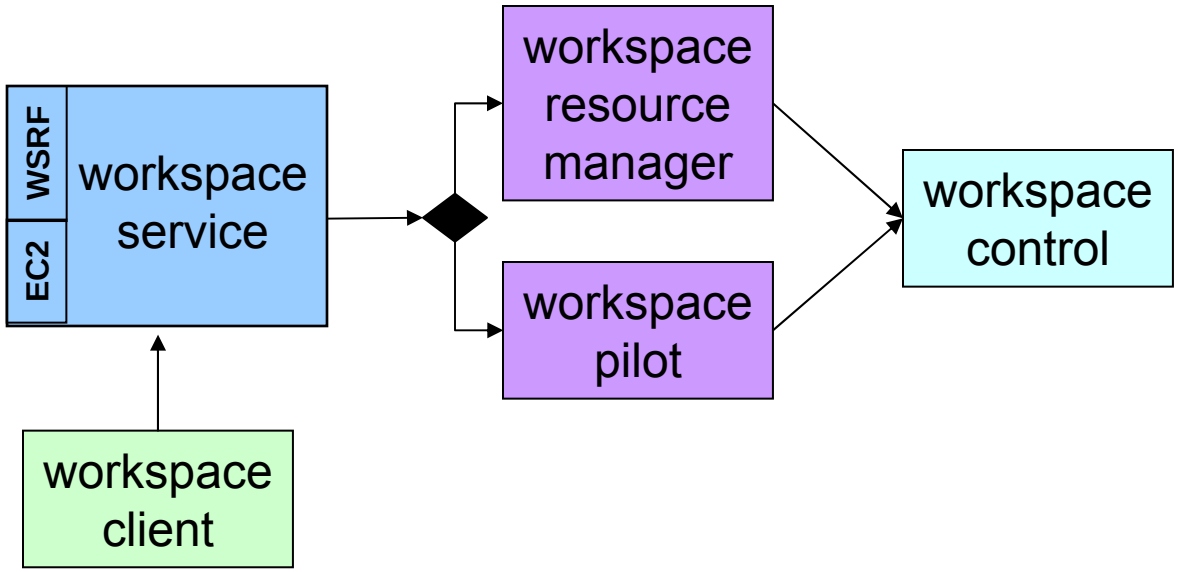

Cloud Closure

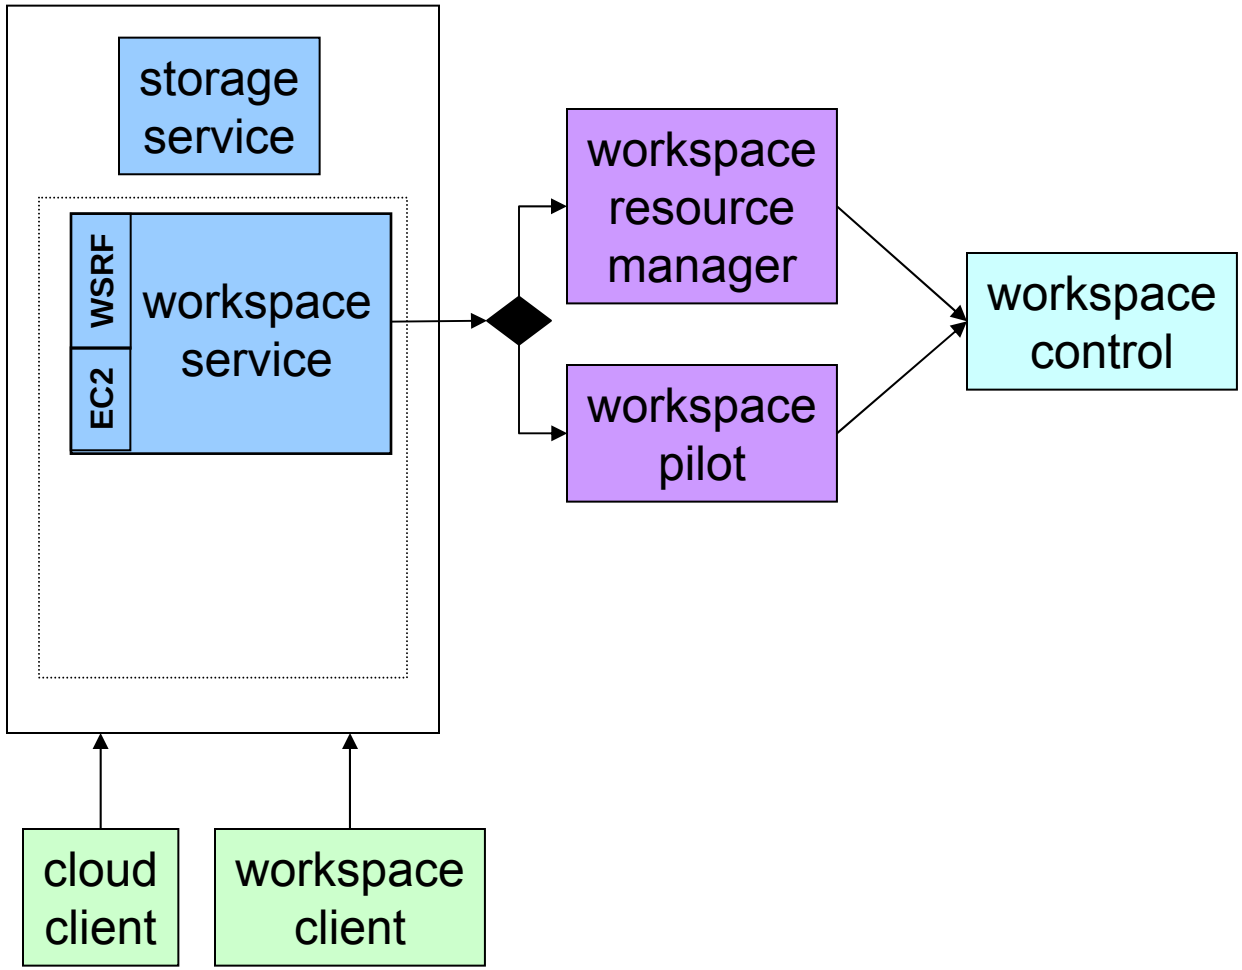

The IaaS Gateway

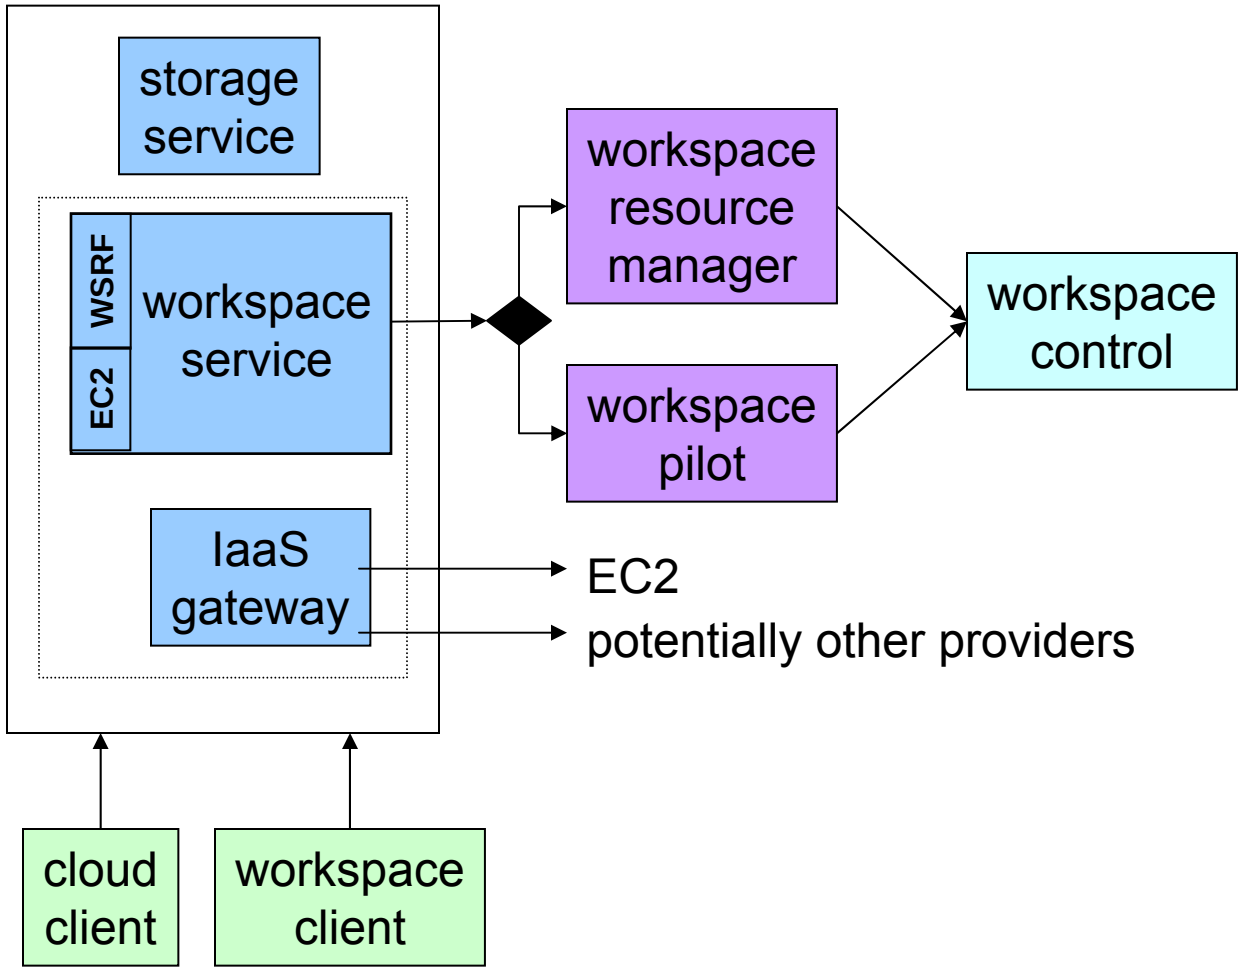

Virtual Clusters

Tightly-coupled clusters

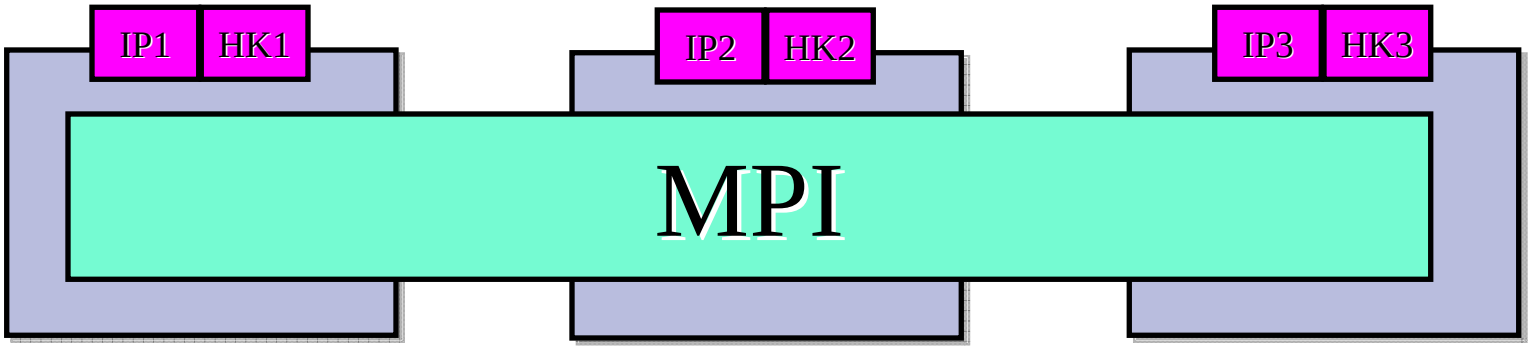

- What makes a cluster a cluster?
  - Shared trust/security context
  - Shared configuration/context information
- Reciprocal exchange of information: networking and security

Context Broker

- Parameterizable appliance
- Context information exchange

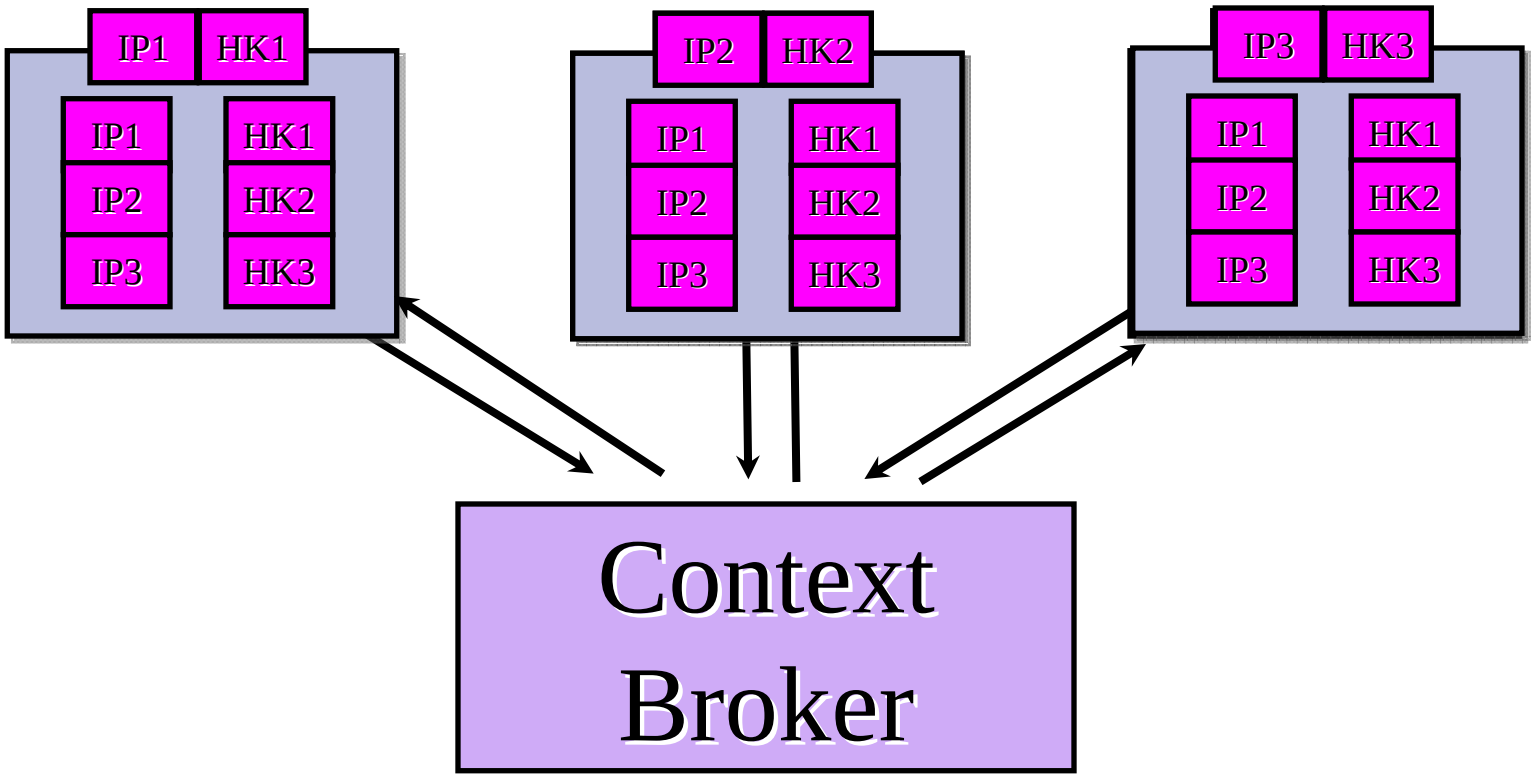

Context Broker Goals

- Can work with every appliance
  - Appliance schema, can be implemented in terms of many configuration systems
- Can work with every cloud provider
  - Simple and minimal conditions on generic context delivery
- Can work across multiple cloud providers, in a distributed environment

Status for Context Broker

- Release history:
  - In alpha testing since August '07
  - First released summer July '08 (v 1.3.3)
  - Latest update January '09 (v 2.2)
- Used to contextualize 100s of nodes for EC2 HEP STAR runs, Hadoop nodes, HEP Alice nodes…
- Contextualized images on workspace marketplace
- Working with rPath to make contextualization easier for the user
  - OVF extensions to be submitted to DMTF

End of Nimbus Tour

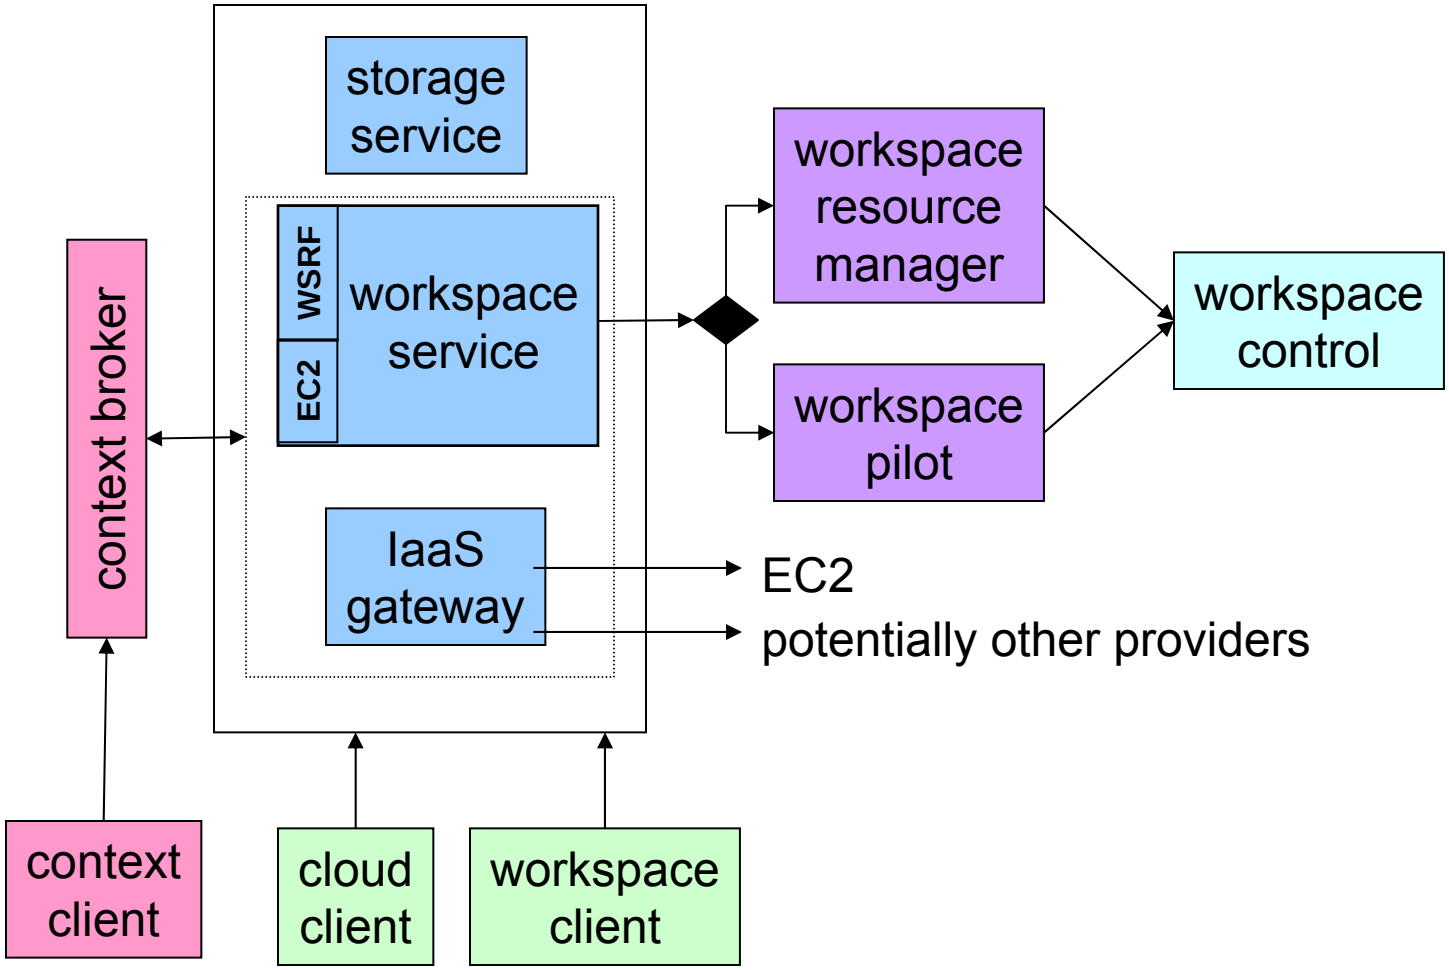

Science Clouds Goals

- Make it easy for scientific projects to experiment with cloud computing
  - Can cloud computing be used for science?
- Evolve 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
- Provide a laboratory for exploration of cloud interoperability issues

Science Cloud Resources

- University of Chicago (Nimbus):
  - first cloud, online since March 4th 2008
  - 16 nodes of UC TeraPort cluster, public IPs
- University of Florida
  - Online since 05/08
  - 16-32 nodes, access via VPN
- Other Science Clouds
  - Masaryk University, Brno, Czech Republic (08/08), Purdue (09/08)
  - Installations in progress: IU, Grid5K, Vrije, others
- Using EC2 for overflow
- Minimal governance model
- http://workspace.globus.org/clouds

Who Runs on Nimbus?

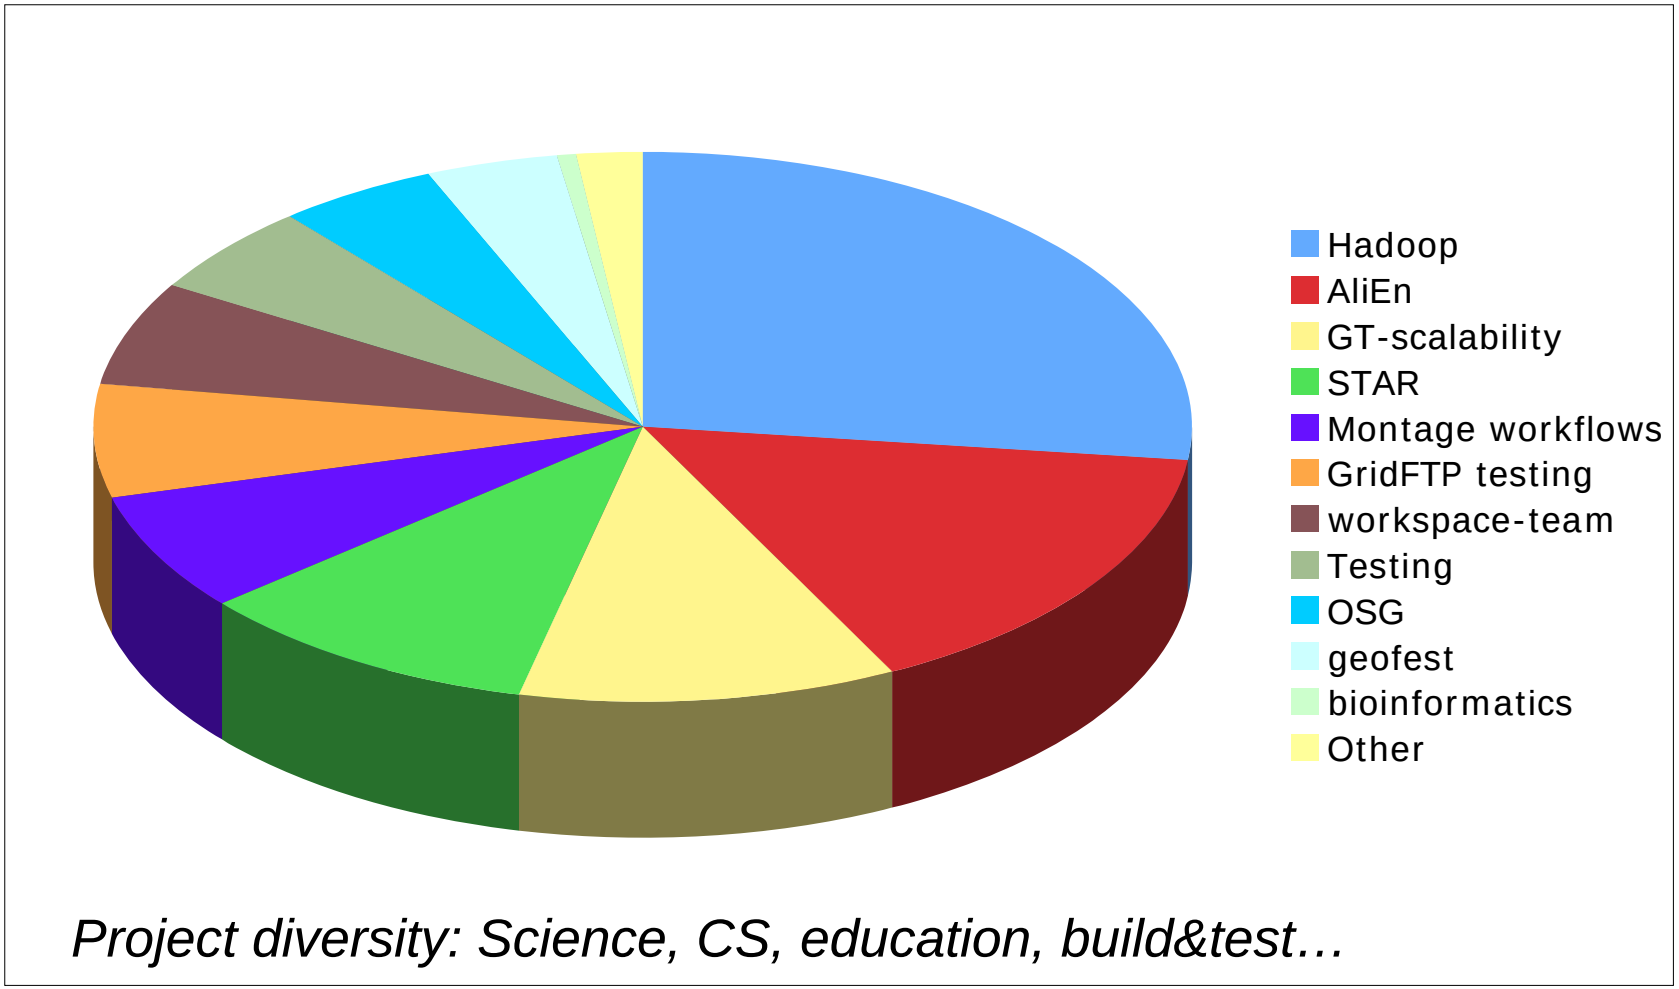

Project diversity: Science, CS, education, build&test…

Hadoop over ManyClouds

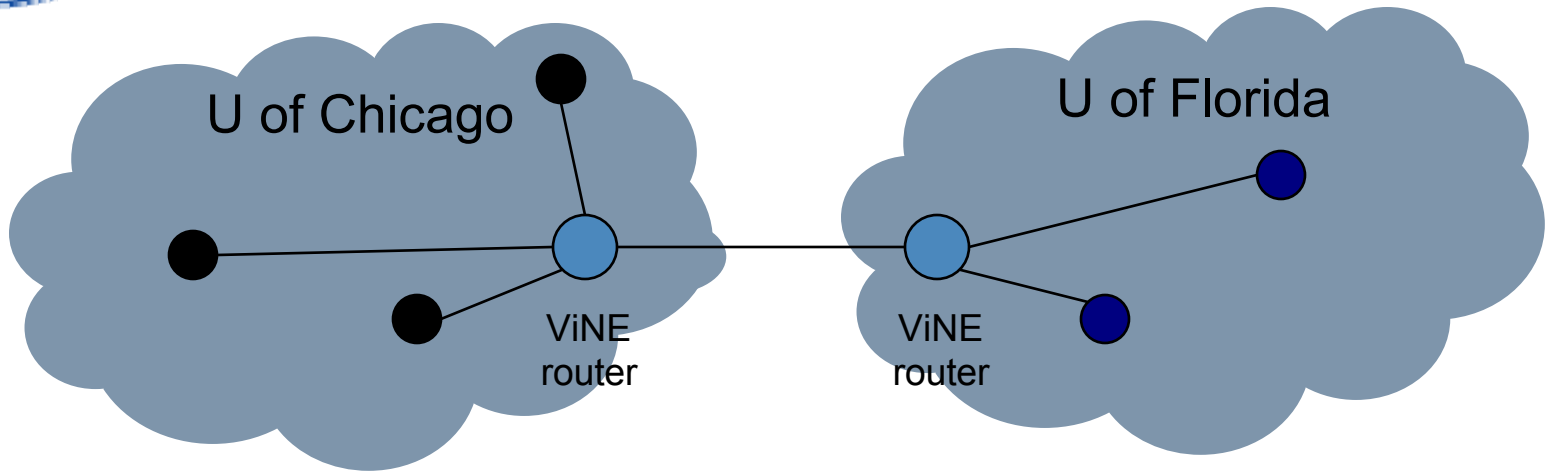

- Building clouds on top of clouds
- Virtual workspace: ViNE router + application VMs
- Need access to distributed resources, and high level of privilege to run a ViNE router
- *Papers:*
  - *"Sky Computing", by K. Keahey, A. Matsunaga, M. Tsugawa, J. Fortes. Submitted to IEEE Internet Computing.*
  - *"CloudBLAST: Combining MapReduce and Virtualization on Distributed Resources for Bioinformatics Applications" by A. Matsunaga, M. Tsugawa and J. Fortes. eScience 2008.*

Alice HEP Experiment at CERN

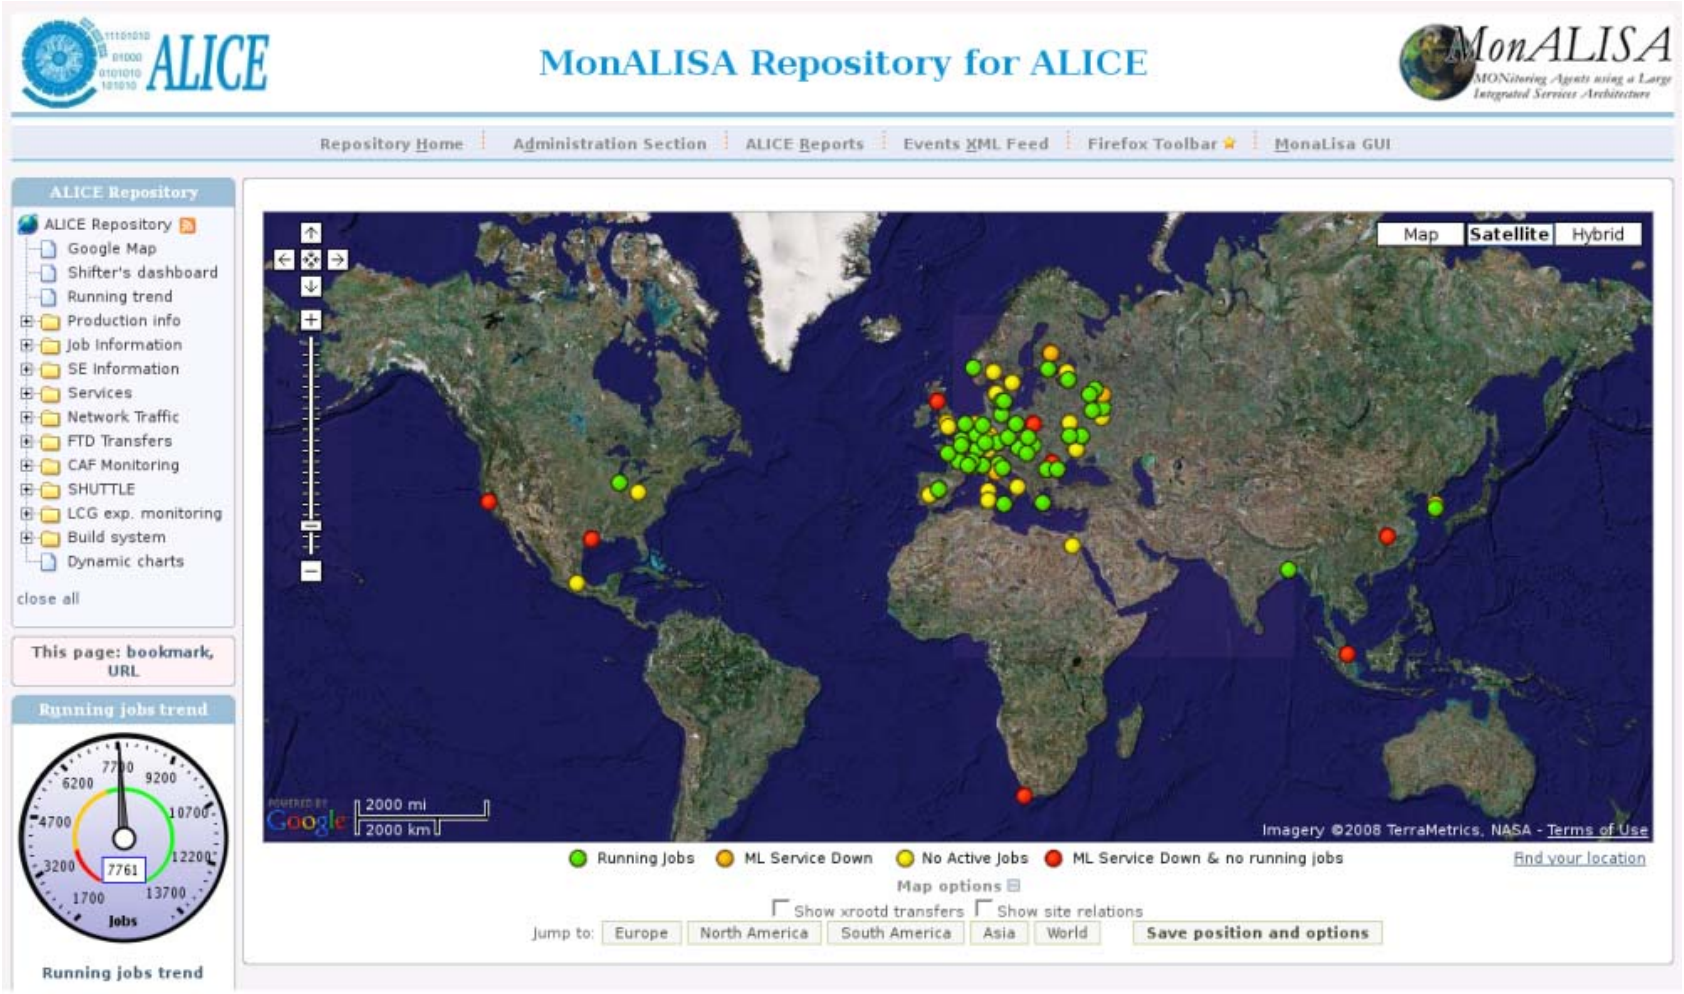

- *Collaboration with CERNVM project (CHEP09 paper)*

STAR

- STAR: a high-energy physics experiment
- Need resources **with the right configuration**
  - Complex environments
  - Consistent environments
- A virtual OSG STAR cluster
  - OSG cluster: OSG CE (headnode), gridmapfiles, host certificates, NSF, Torque, worker nodes: SL4 + STAR
- Requirements
  - One-click virtual cluster deployment
  - Science Clouds -> EC2
- From proof-of-concept to productions runs
- Work by Jerome Lauret, Doug Olson, Leve Hajdu, Lidia Didenko
  - Results to be published at Quark Matter conference and CHEP

Infrastructure Testing

- Motivation
  - Test middleware scalability, use of different platforms, etc.
- Workspaces
  - Globus 101 and several different environments
- Requirements
  - very short-term but flexible access to diverse platforms
- Work by various members of the Globus community (Tom Howe and John Bresnahan), short-lived "communities of one"
- Resulted in provisioning a private cloud for Globus

Montage Workflows

- 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*

Friends and Family

- Committers: Kate Keahey & Tim Freeman (ANL/UC), Ian Gable (UVIC)
- A lot of help from the community, see: http://workspace.globus.org/people.html
- Collaborations:
  - Cumulus: S3 implementation (Globus team)
  - EBS implementation with IU
  - Appliance management: rPath and Bcfg2 projects
  - Virtual network overlays: University of Florida
  - Security: Vienna University of Technology

Nimbus Users

- Applications users
  - Scientific projects
  - Use clouds, Nimbus client side, user guides
- Cloud administrators
  - Resource providers
  - Install Nimbus, administrator guides
- Communities extending Nimbus
  - Extensions for research or usability
  - Develop code, extensibility guides

Open Source IaaS Implementations

- OpenNebula
  - Open source datacenter implementation
  - University of Madrid, I. Llorente & team, 03/2008
- Eucalyptus
  - Open source implementation of EC2
  - UCSB, R. Wolski & team, 06/2008
- Cloud-enabled Nimrod-G
  - Open source implementation of EC2
  - Monash University, MeSsAGE Lab, 01/2009
- Industry efforts
  - openQRM, Enomalism

Cloud Computing Ecosystem

Appliance Providers

marketplaces
commercial providers
communities (CERNVM)

Deployment Orchestrator

orchestrate the deployment of environments across possibly many cloud providers

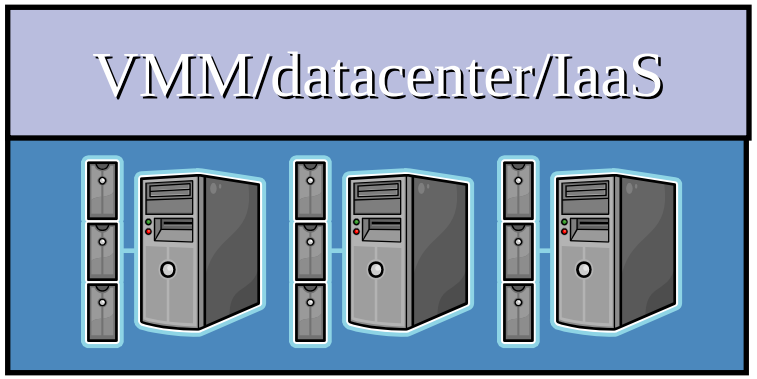

IaaS Clouds vs Grids

- Interfaces to resources in the cloud
- Significant new abstraction: provisioning resources and environments/configurations
  - Configuration control
  - Resource control, better articulated SLAs
- Revolution or evolution?
  - We can leverage much of the infrastructure developed in the context of grid computing
  - There is new potential to explore

the globus alliance
www.globus.org

Some Thoughts

- Cloud computing has an enormous potential for Science
- New roles: appliance providers
  - Or maybe not so new...
- Interoperability
  - Academic vs commercial resources
  - Standards: “rough consensus & working code”
- Importance of open source
  - Drive academic needs into the marketplace
  - Drive the development of standards
- End-2-end tools
  - Combine with what we have
  - Explore new potential

Acknowledgements

- NSF CSR “Virtual Playgrounds”
- NSF SDCI “Missing Links”
- DOE CEDPS proposal
- Various other collaborations with TeraGrid, OOI, and other communities