Nimbus: Open Source Infrastructure-as-a-Service Cloud Computing Software

Workshop on adapting applications and computing services to multi-core and virtualization

CERN, June 2009

Kate Keahey

keahey@mcs.anl.gov

Nimbus project lead

University of Chicago

Argonne National Laboratory

Cloud Computing for Science

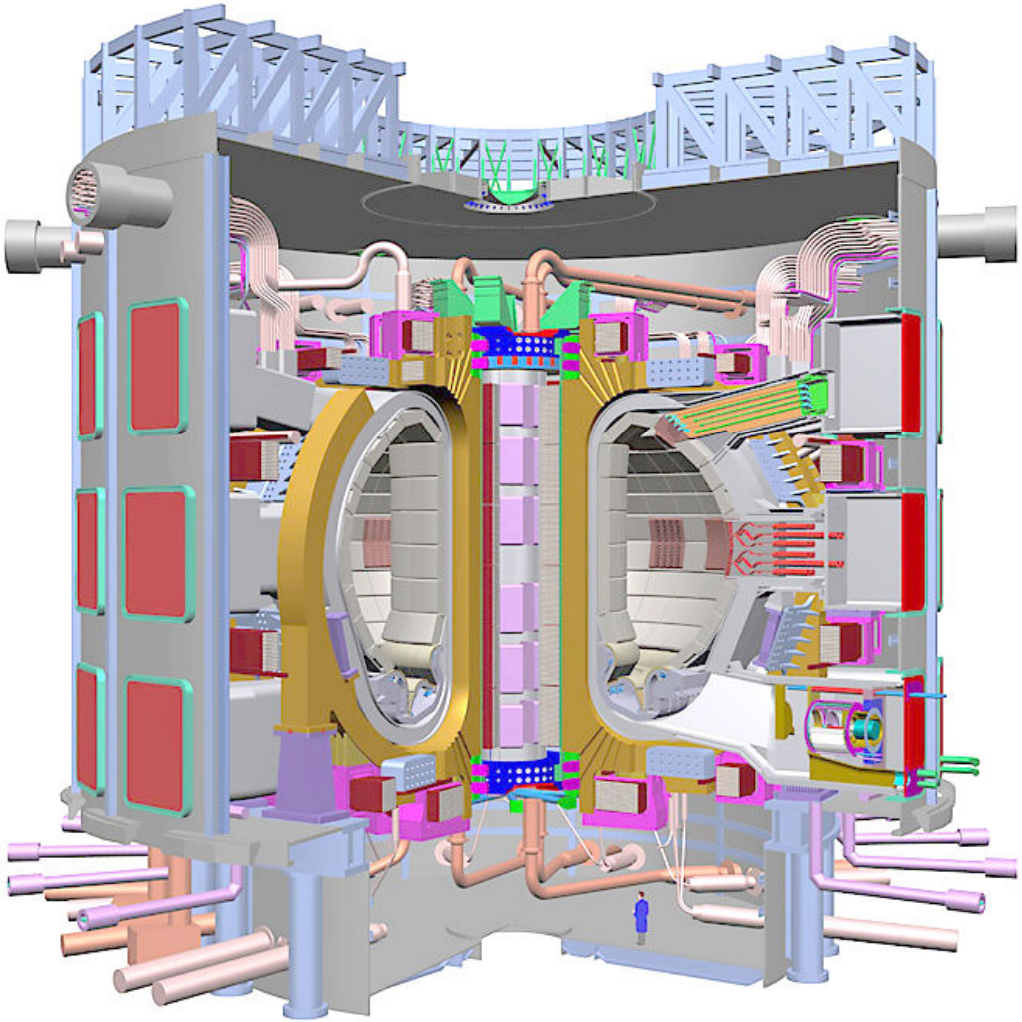

- Complex environments
- Need for control

Workspaces

- Dynamically provisioned environments
  - Environment control
  - Resource control
- Implementations
  - Via leasi… …s: reimaging, confi… …mic accou…
  - Via virtua… …oyment

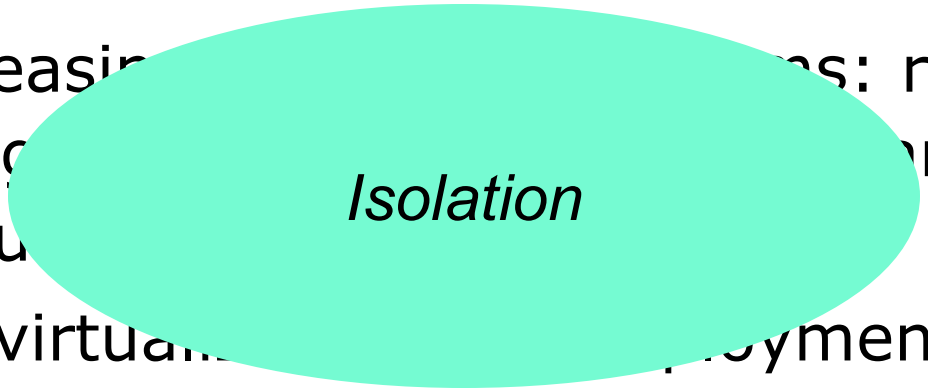

A Brief History of Nimbus

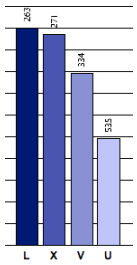

Xen released

"A Case for Grid Computing on VMs"
In-Vigo, VIOLIN, DVEs, Dynamic accounts
Policy-driven negotiation

EC2 goes online

First WSRF Workspace Service release

STAR ☆

First STAR production run on EC2

EC2 gateway available

Nimbus Cloud comes online

Context Broker release

Support for EC2 interfaces

2003 2006 2009

Nimbus Goals

- Allow providers to build clouds
  - Private clouds (privacy, expense considerations)
  - Workspace Service: open source EC2 implementation
- Allow users to use cloud computing
  - Do whatever it takes to enable scientists to use IaaS
  - Context Broker: turnkey virtual clusters
  - IaaS Gateway: interoperability
- Allow developers to experiment with Nimbus
  - For research or usability/performance improvements
  - Community extensions and contributions

The Workspace Service

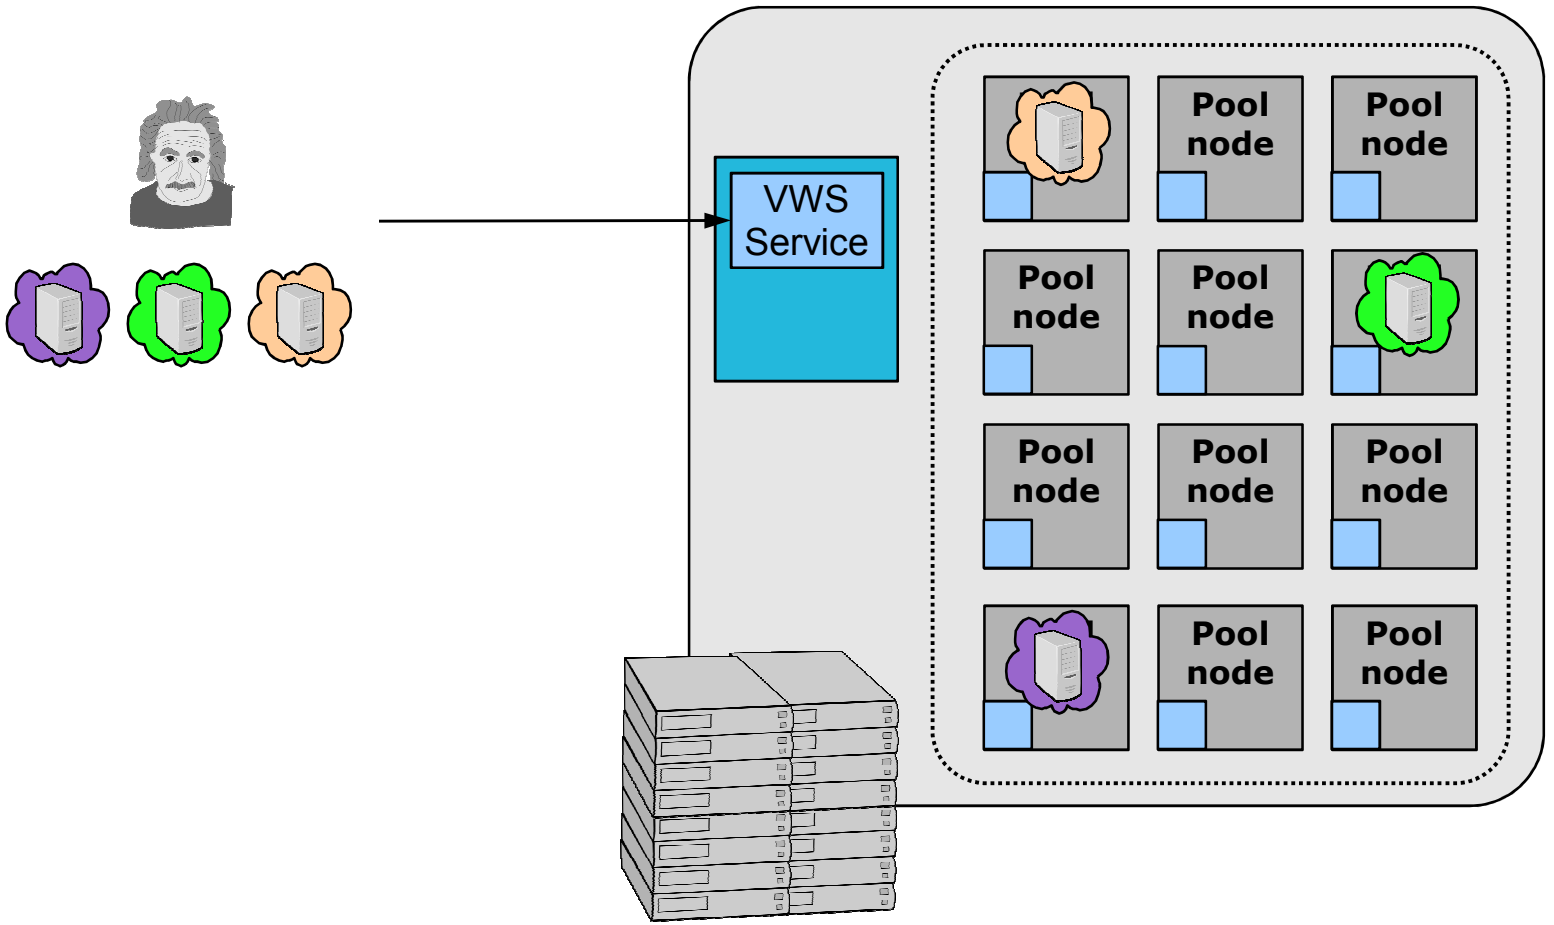

The Workspace Service

The workspace service publishes information about each workspace

Users can find out information 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.

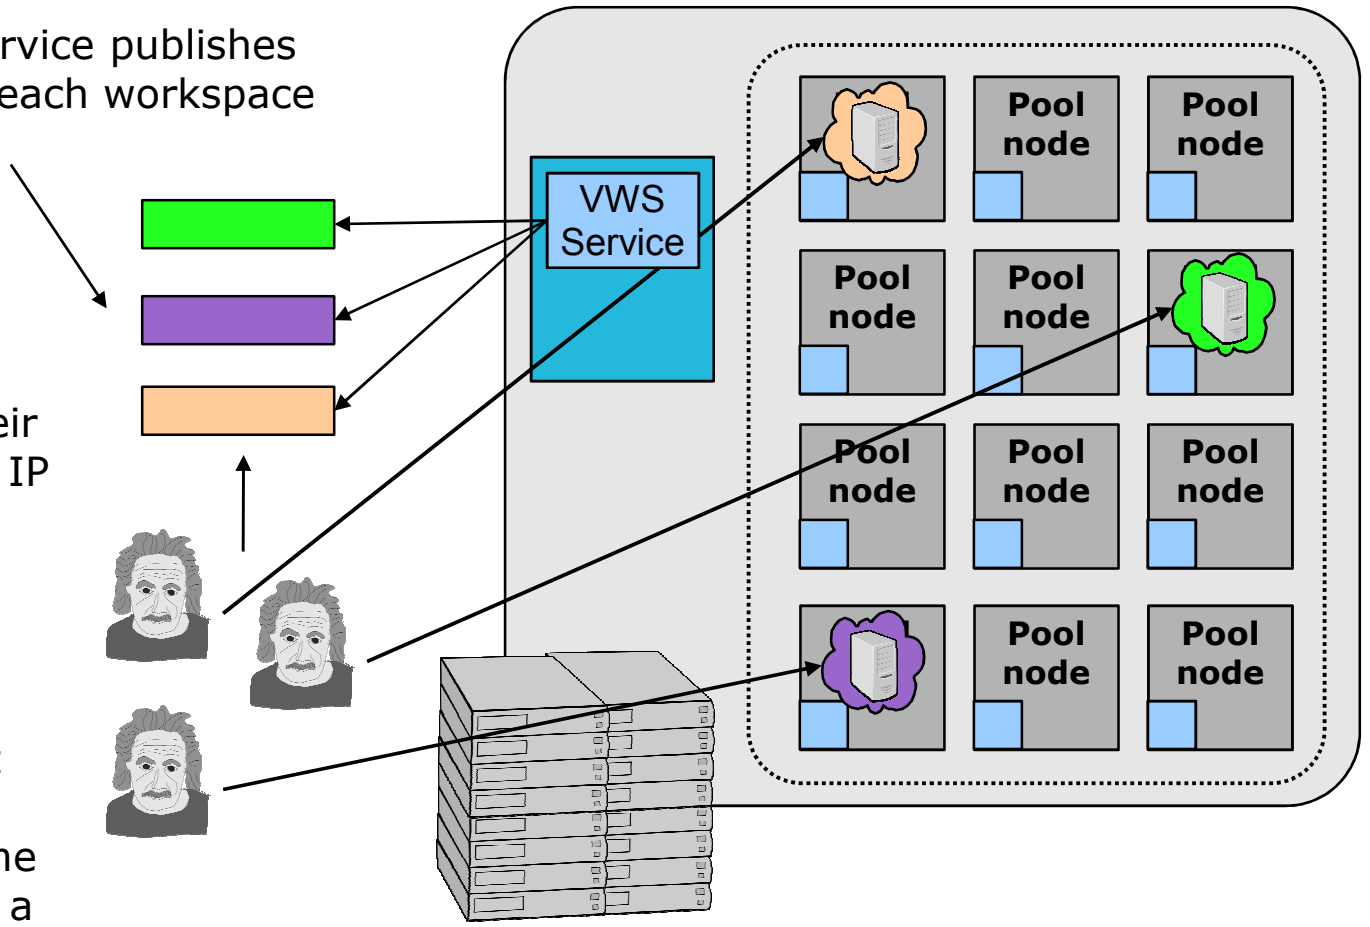

Workspace Service: Interfaces and Clients

- Web Services based
- Web Service Resource Framework (WSRF)
  - WS + state management (WS-Notification)
- Elastic Computing Cloud (EC2)
  - Compatible with EC2 clients
  - 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
- Protocol adapter, moving towards messaging

Workspace Service: Security

- GSI authentication and authorization
  - PKI-based
  - VOMS, Shibboleth (via GridShib), custom PDPs
- Secure access to VMs
  - EC2 key generation or accessed from .ssh
- Validating images and image data
  - Extensions from Vienna University of Technology
  - *Paper: Descher et al., Retaining Data Control in Infrastructure Clouds, ARES (the International Dependability Conference), 2009.*

Workspace Service: 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

Workspace Components

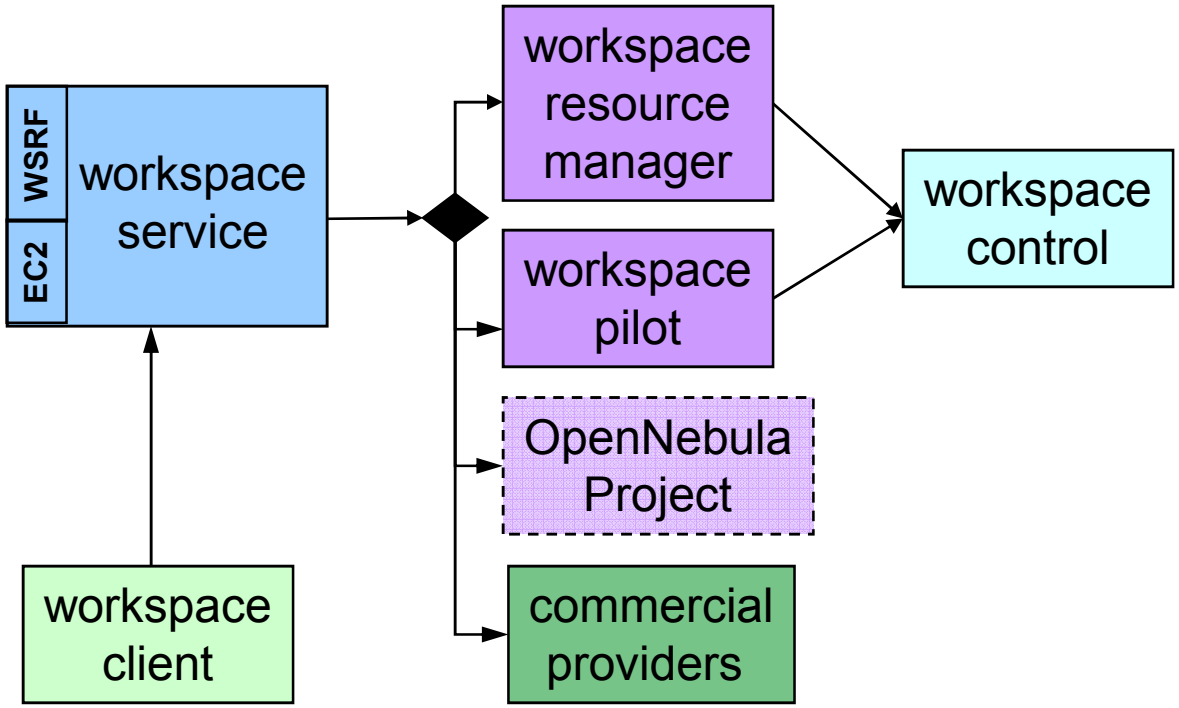

See papers at: http://workspace.globus.org/papers/index.html
- "Simple Leases with Workspace Pilot" (EuroPar08)
- "Combining Batch Execution and Leasing Using Virtual Machines" (HPDC08),

Cloud Capabilities

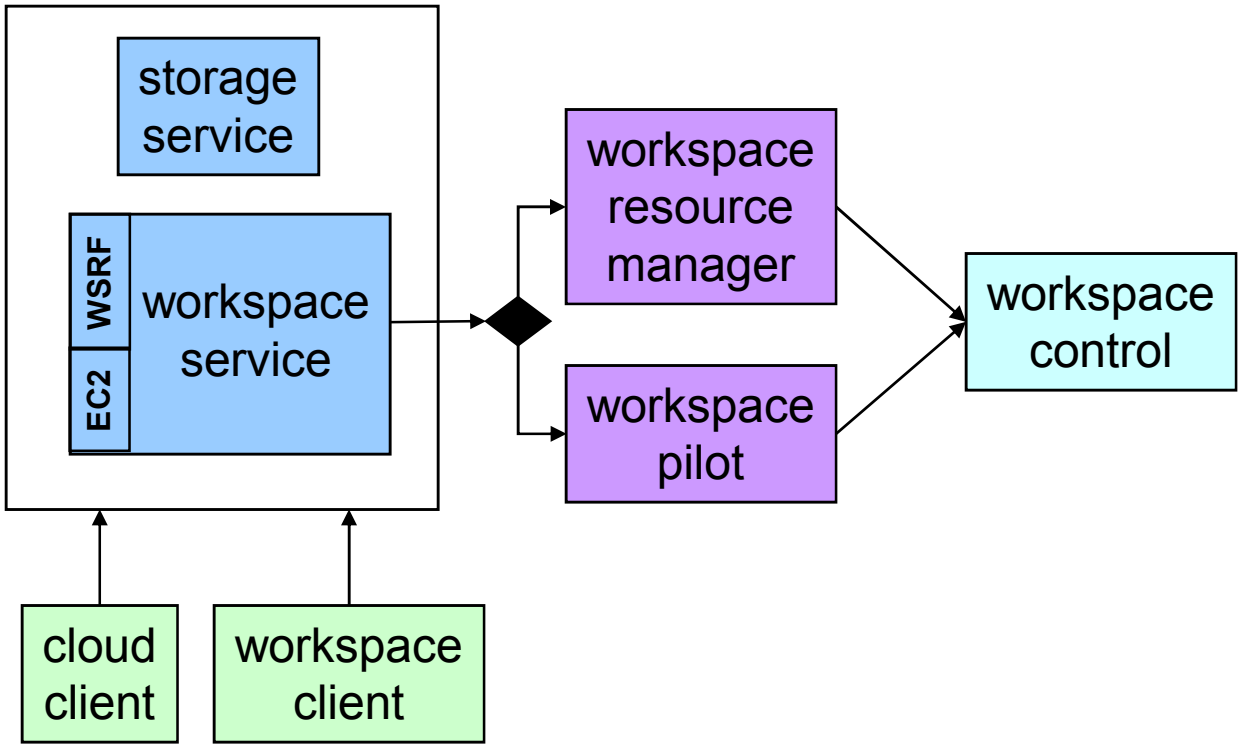

Cloud Computing Ecosystem

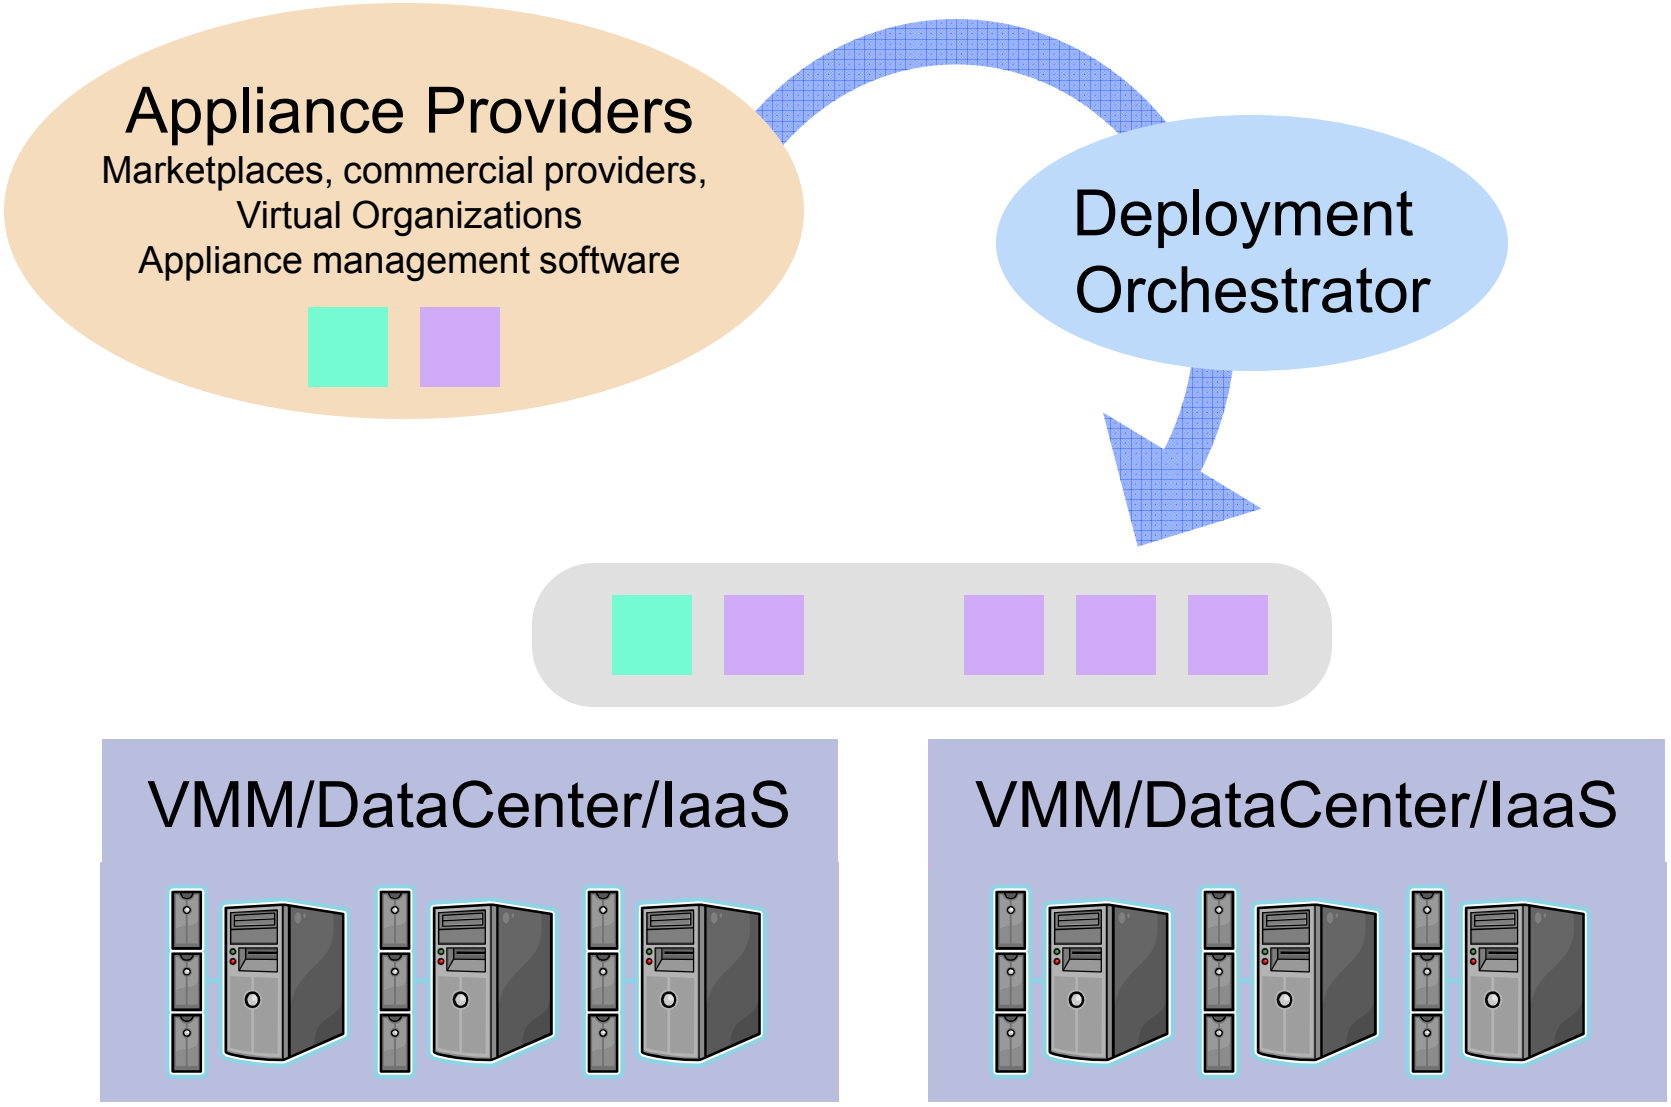

Turnkey Virtual Clusters

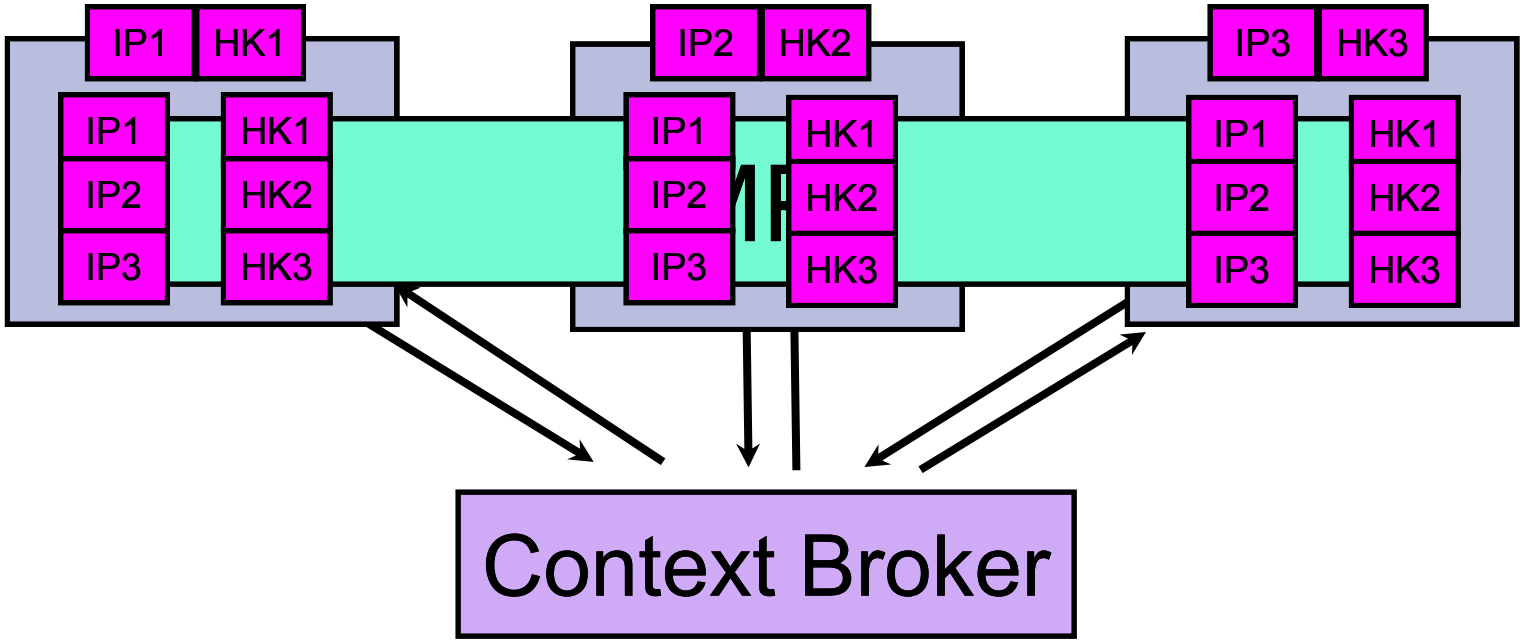

- Turnkey, tightly-coupled cluster
  - Shared trust/security context
  - Shared configuration/context information

Context Broker Status

- Releases
  - In alpha since 08/07, first release 06/08, update 01/09
- Used to contextualize cluster composed of 100s of virtual nodes for multiple production apps
- Contextualized images on workspace marketplace
- Working with rPath to make contextualization easier for the user

Paper: Keahey&Freeman, Contextualization: Providing One-Click Virtual Clusters, eScience 2008

End of Nimbus Tour

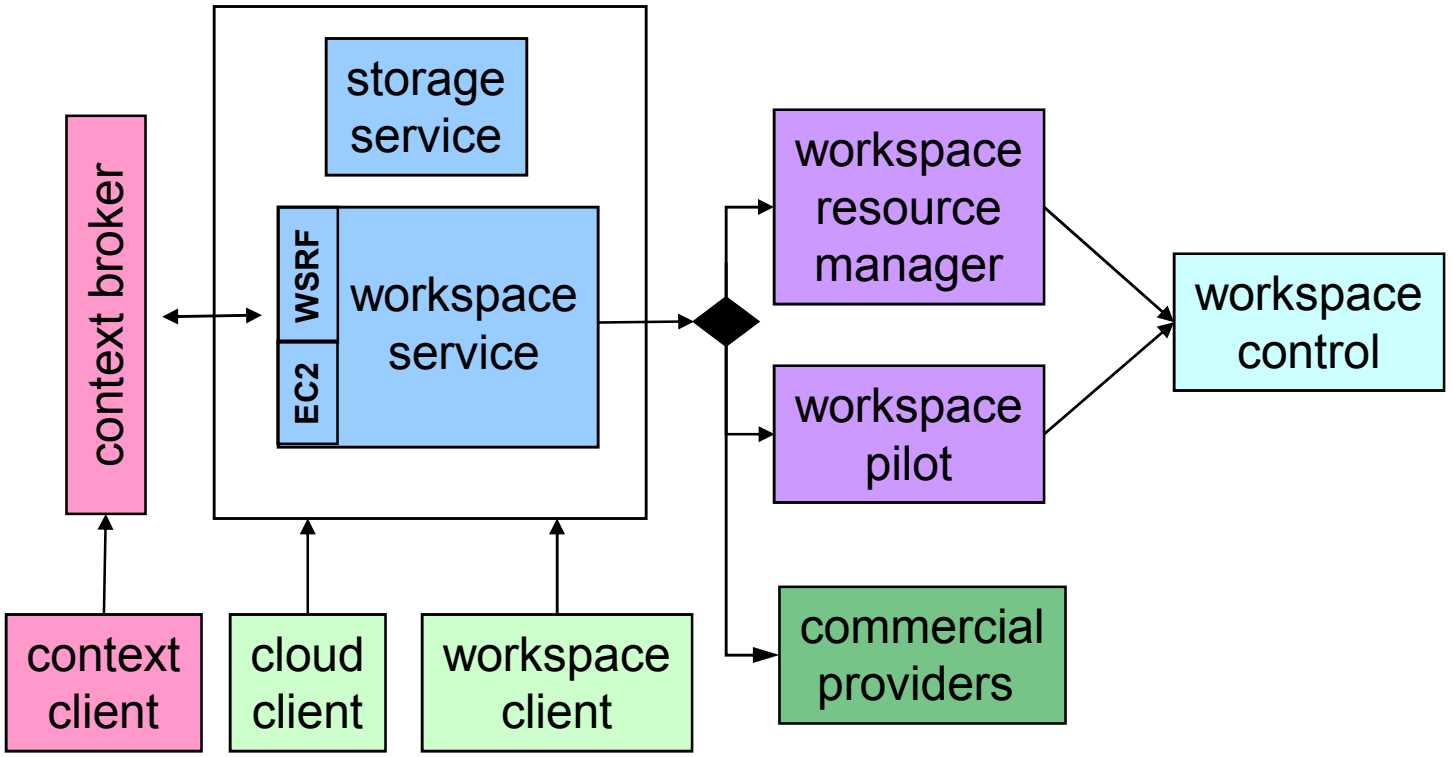

Nimbus: Extensions and Collaborations

- Nimbus core team:
  - UC/ANL: Kate Keahey, Tim Freeman, David LaBissoniere
  - UVIC: Ian Gable & team (Nimbus monitoring)
- Cumulus: Raj Kettimuthu and John Bresnahan (ANL/UC)
- EBS: Marlon Pierce, Xiaoming Gao, Mike Lowe (IU)
- ViNe: Mauricio Tsugawa, Jose Fortes (UFL)
- Others:
  - OpenNebula project (University of Madrid)
  - Descher et al (Technical U of Vienna): privacy extensions

Congratulations to the EBS team for winning best poster award at TG09!

Science Clouds

- Goals
  - Enable experimentation with IaaS
  - Evolve software in response to user needs
  - Exploration of cloud interoperability issues
- Participants
  - University of Chicago (since 03/08, 16 nodes), University of Florida (05/08, 16-32 nodes, access via VPN), Masaryk University, Brno, Czech Republic (08/08), Wispy @ Purdue (09/08)
  - In progress: Grid5K, IU, Vrije
  - Using EC2 for large runs
- Science Clouds Marketplace: OSG cluster, Hadoop, etc.
- Come and run: http://workspace.globus.org/clouds

Who Runs on Nimbus?

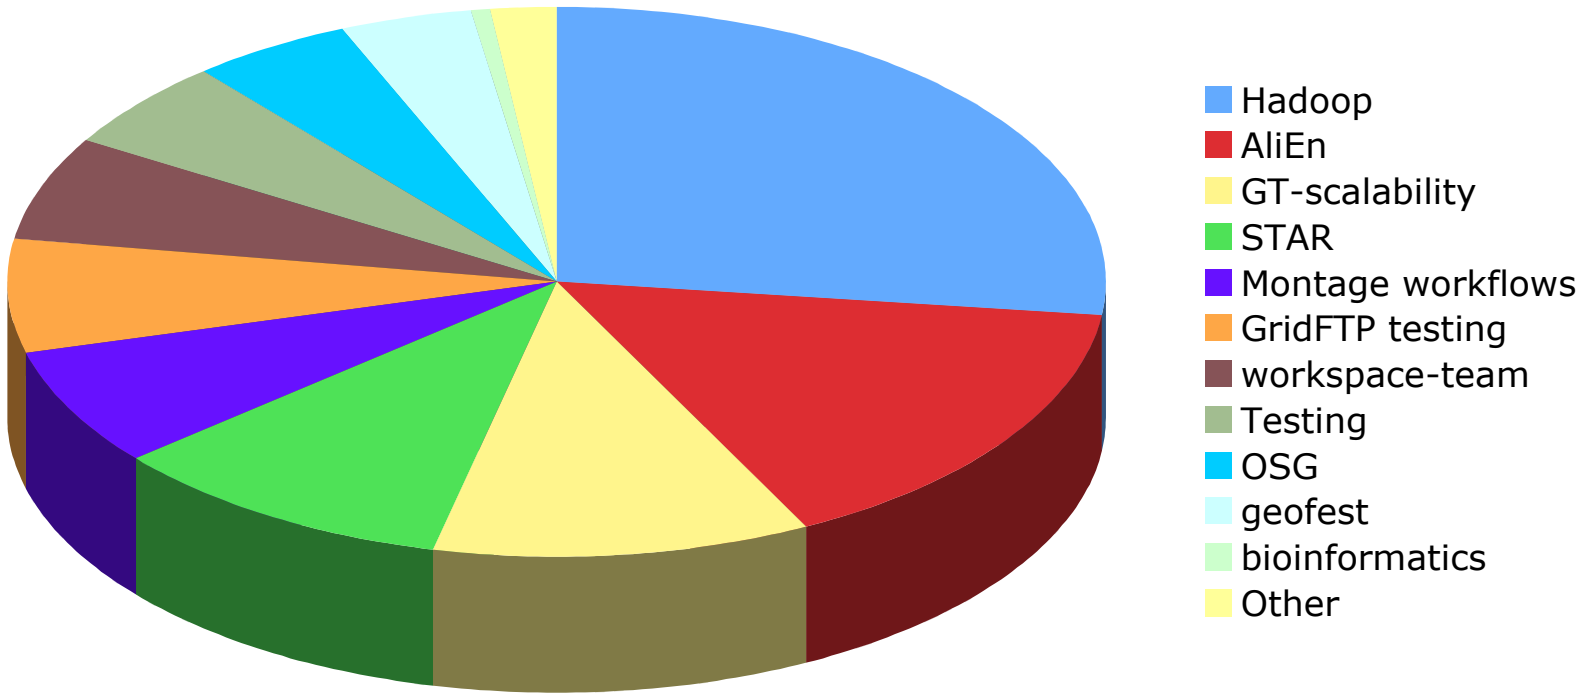

100+ DNs

projects ranging across Science, CS, education, build&test…

STAR experiment

- STAR: a nuclear physics experiment at Brookhaven National Laboratory
- Studies fundamental properties of nuclear matter
- Problem: computations require complex and consistently configured environments that are hard to find in existing grids

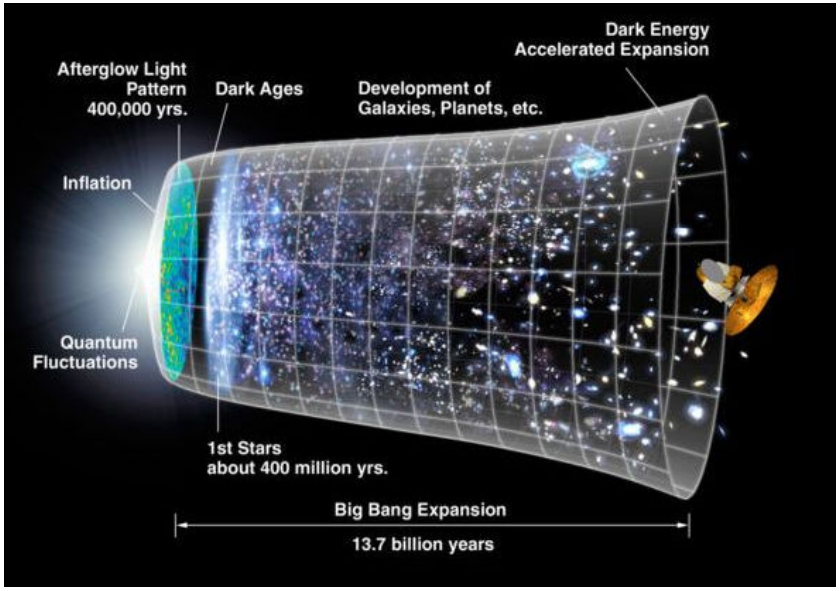

STAR Virtual Clusters

- Virtual resources
  - A virtual OSG STAR cluster: OSG headnode (gridmapfiles, host certificates, NFS, Torque), worker nodes: SL4 + STAR
  - One-click virtual cluster deployment via Nimbus Context Broker
- From Science Clouds to EC2 runs
- Running production codes since 2007
- Work by Jerome Lauret, Leve Hajdu, Lidia Didenko (BNL), Doug Olson (LBNL)
- The Quark Matter run: producing just-in-time results for a conference: http://www.isgtw.org/?pid=1001735

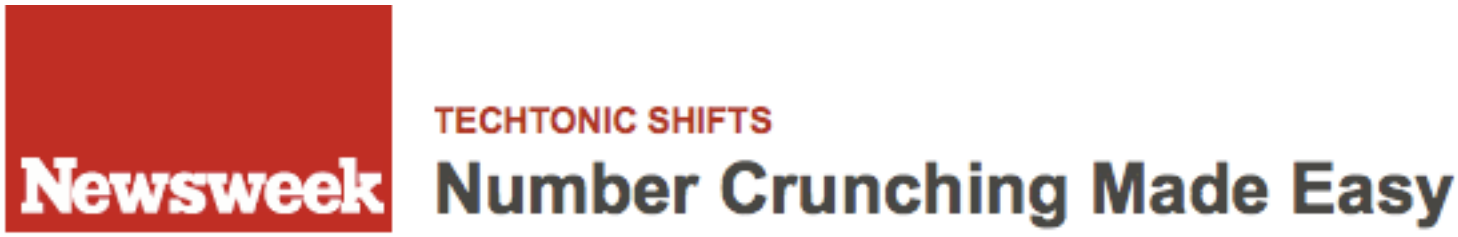

STAR Quark Matter Run

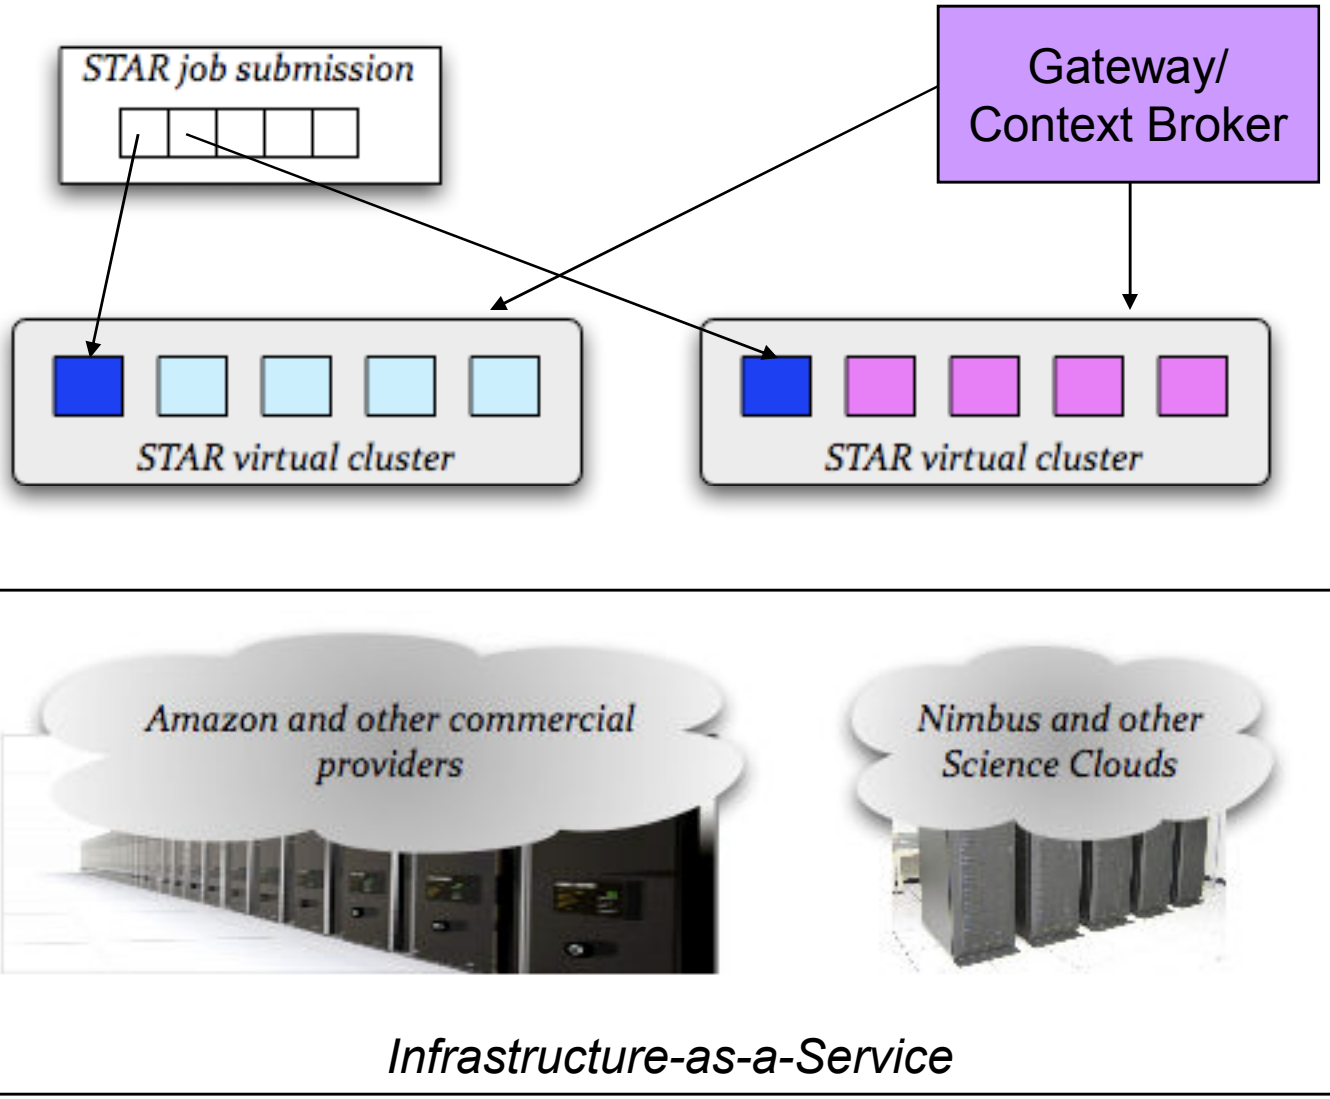

STAR Quark Matter Run (2)

- Application stats:
  - Processed 1.2 M events
  - Moved ~1TB of data over duration (small I/O needs)
- Run facts:
  - 300+ nodes, total of ~36,000 hours
  - Instances, 32-bit, 1.7 GB memory:
    - EC2 default: 1 EC2 CPU unit
    - High-CPU Medium Instances: 5 EC2 CPU units (2 cores)
  - Cost:
    - Comp: ~ $6,000: ~ $1,7 K (default) + ~ $3,9K (medium)
    - Data: ~ $150

Elastic Provisioning for AliEn

CernVM
SoFTware Appliance

- Challenge: integrate elastic computing into existing infrastructure for the ALICE experiment
- Collaboration with CernVM project
- Work by Artem Harutyunyan

Elastic Provisioning for AliEn

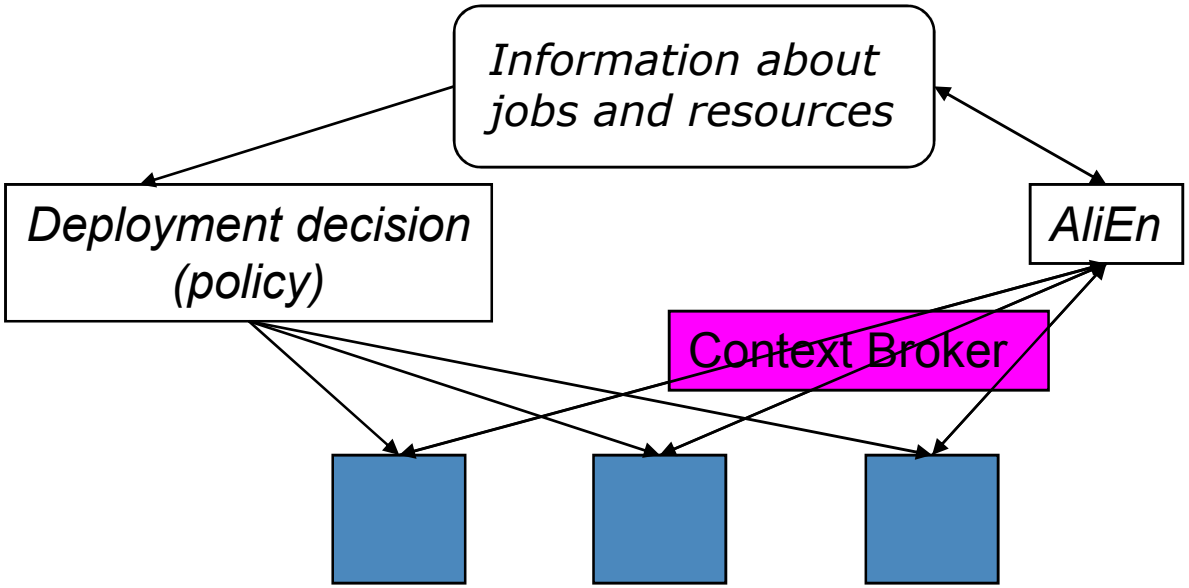

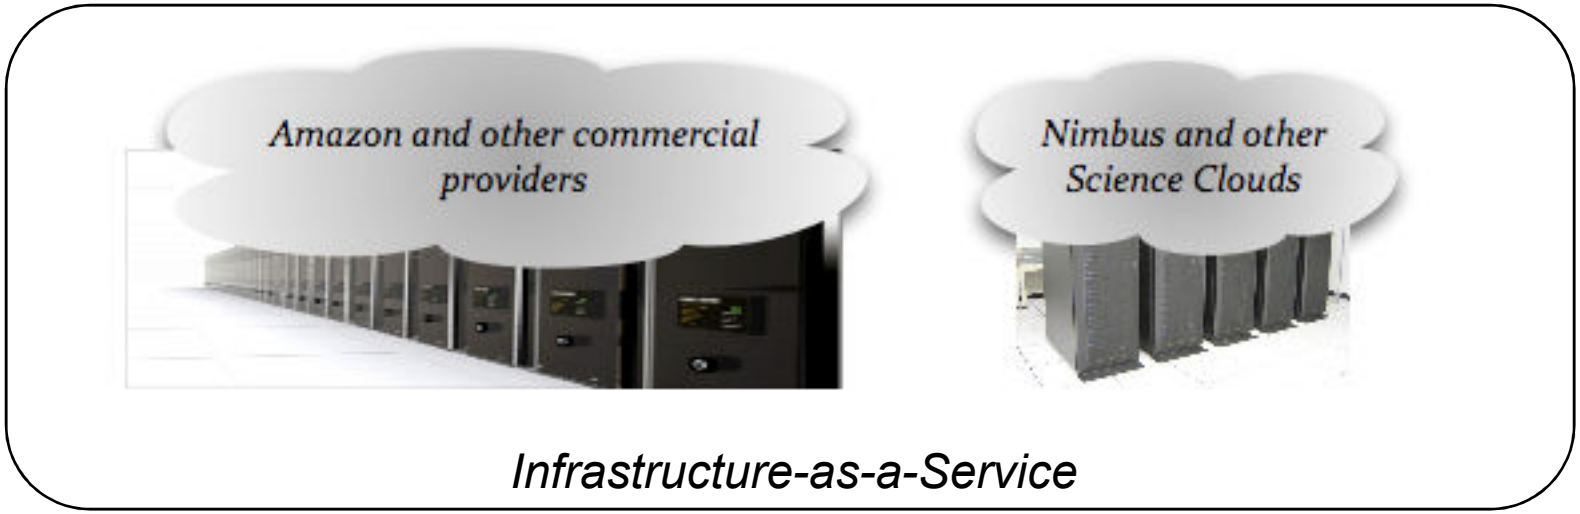

Elastic Provisioning for AliEn

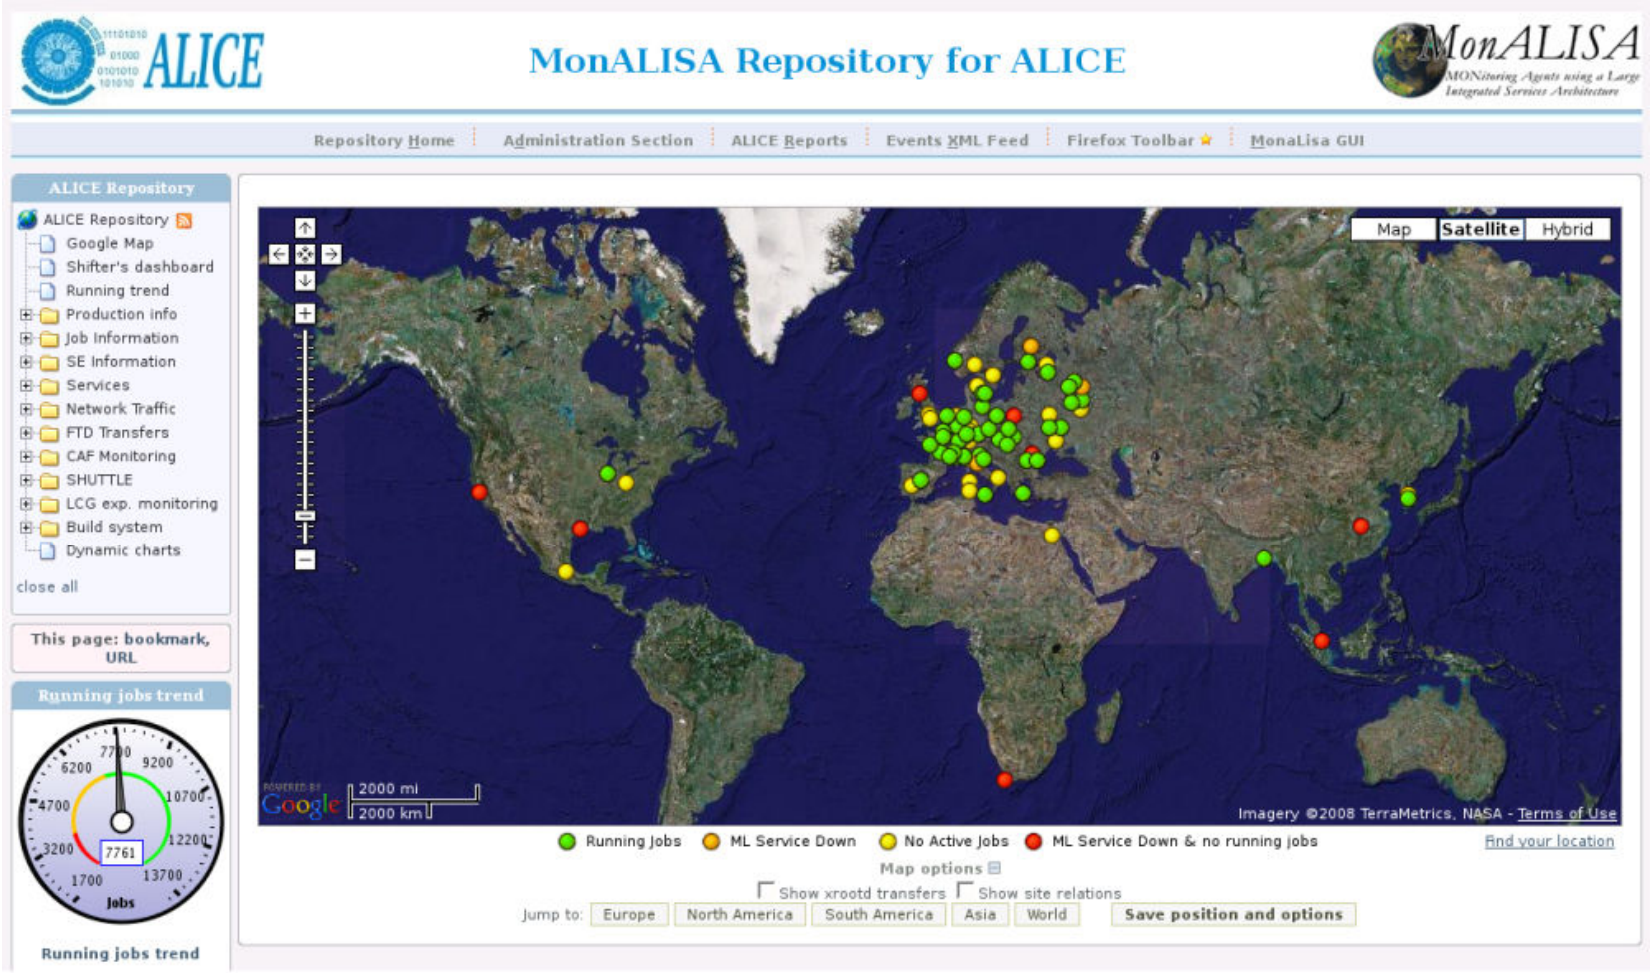

- *CHEP09 paper, Harutyunyan et al.*
- *Similar ongoing work on extensions to GANGA (ATLAS) and extensions to local schedulers*

Sky Computing Environment

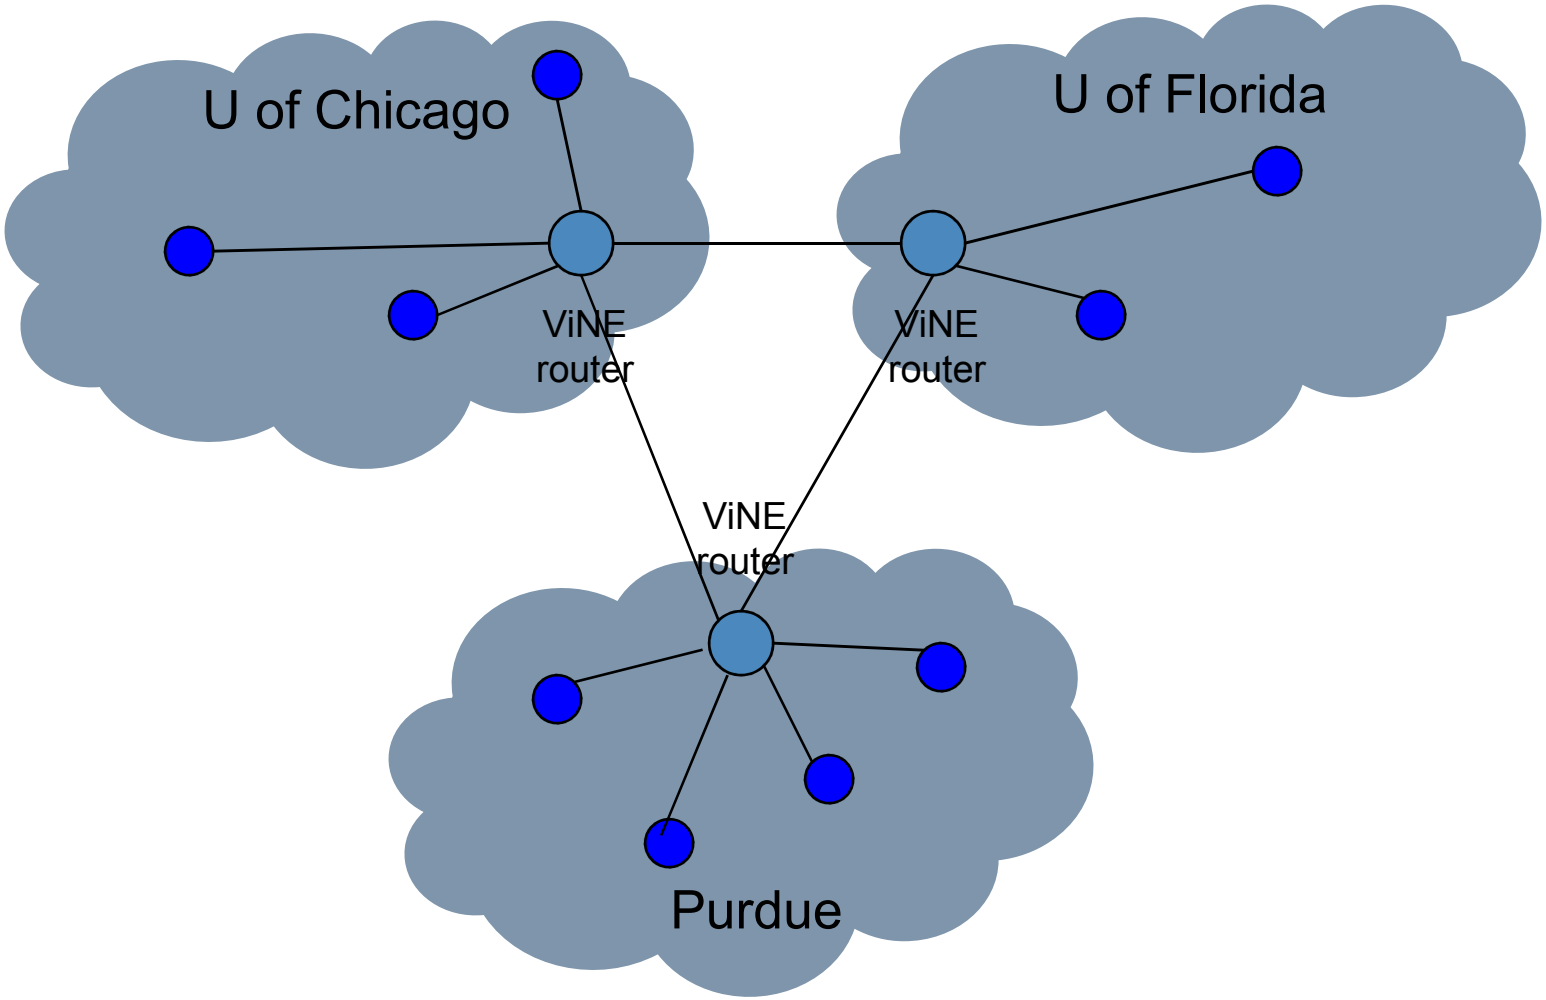

Work by A. Matsunaga, M. Tsugawa, University of Florida

Hadoop in the Science Clouds

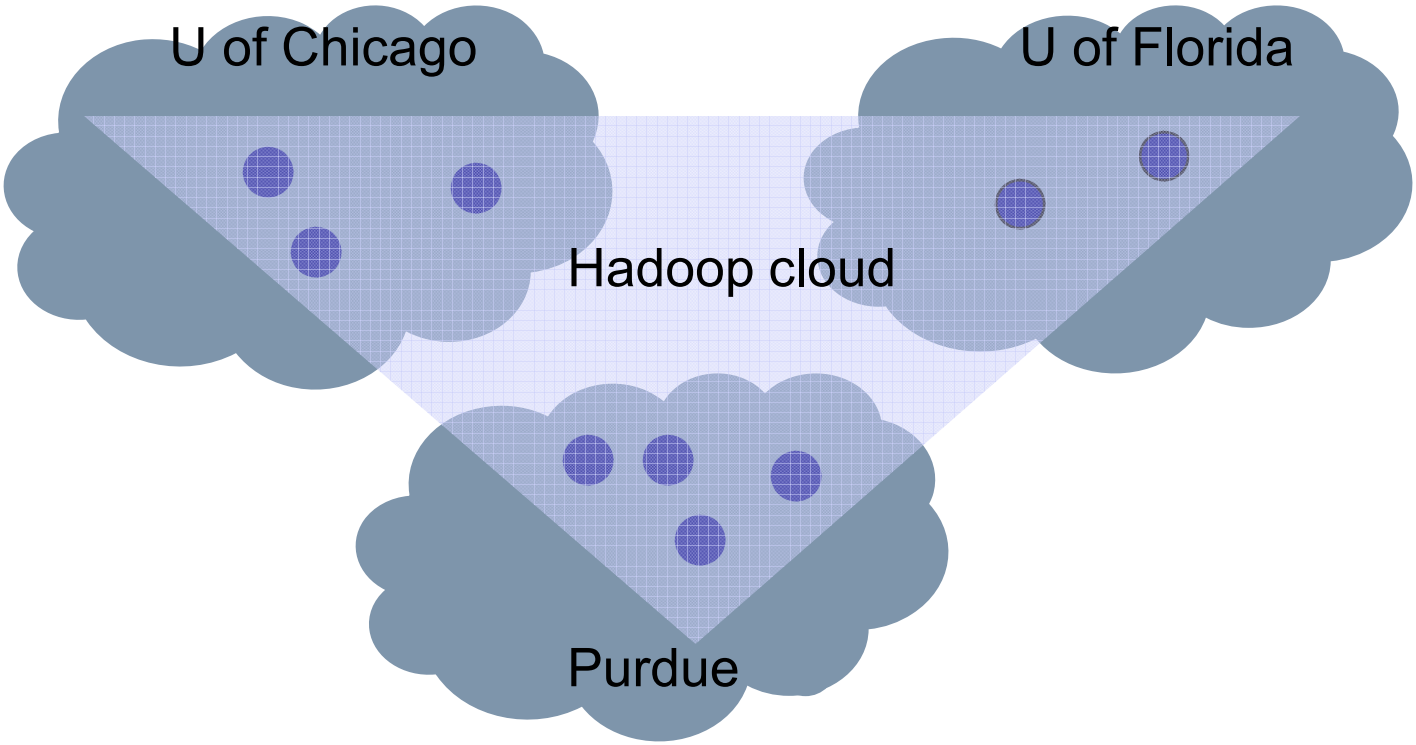

- *Papers:*
  - *"Sky Computing", by K. Keahey, A. Matsunaga, M. Tsugawa, J. Fortes. Submitted to IEEE Internet Computing.*
  - *"CloudBLAST: Combining MapReduce and Virtualization on Distributed Resources for Bioinformatics Applications" by A. Matsunaga, M. Tsugawa and J. Fortes. eScience 2008.*

Parting Thoughts

- IaaS cloud computing is science-driven
- Scientific applications are successfully using the existing infrastructure for production runs
- Many more could be using it, but challenges exist…
- Project for the next few years: solve them!