2nd International ICST Conference on Cloud Computing
(CloudComp 2010)
Barcelona, Spain October 26 – 28, 2010

Cloud versus Cloud: the Blessings and Challenges of Cloud Computing for Science

Kate Keahey
keahey@mcs.anl.gov
Argonne National Laboratory
Computation Institute, University of Chicago

Cloud versus Cloud

You mean I have to install my own OS?

How long does it take to deploy 5000 nodes on a cloud? And how long will the comp take?

This resource is unreliable if it shows up at all! And what about security?

It costs too much! And what if Amazon raises prices?

Custom user environments!

On-demand access!

Elastic computing!

Growth and cost management!

Capital expense -> operational expense!

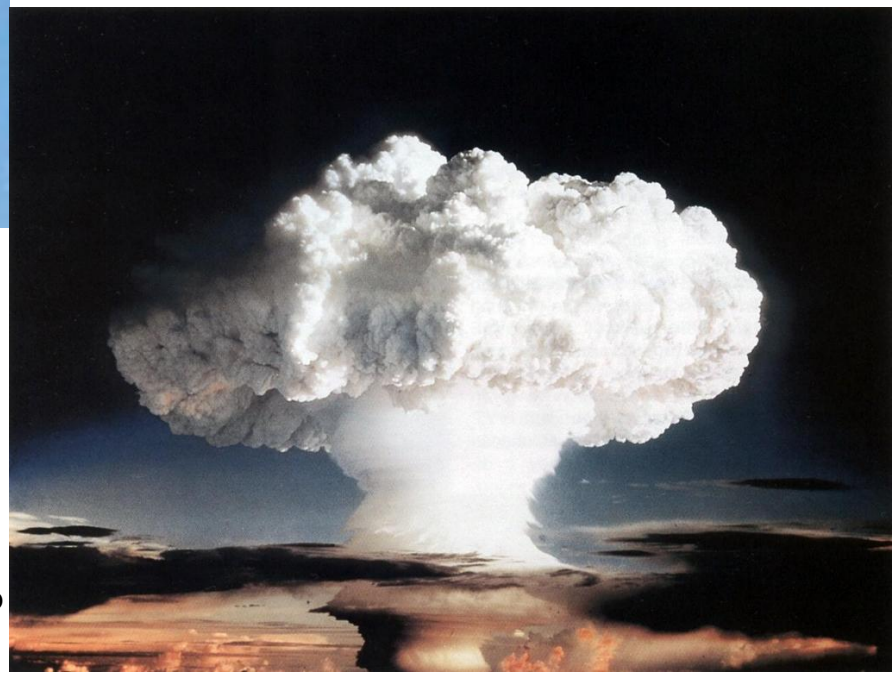

What We Do on the Nimbus Team…

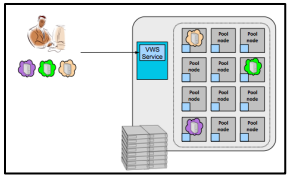

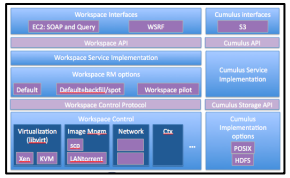

Infrastructure-as-a-Service
Cloud Computing with Nimbus

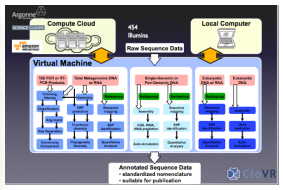

Applications and Ecosystem

Count our Blessings and
Resolve our Challenges

Nuts and Bolts: Infrastructure-as-a-Service Cloud Computing with Nimbus

Nimbus Goals

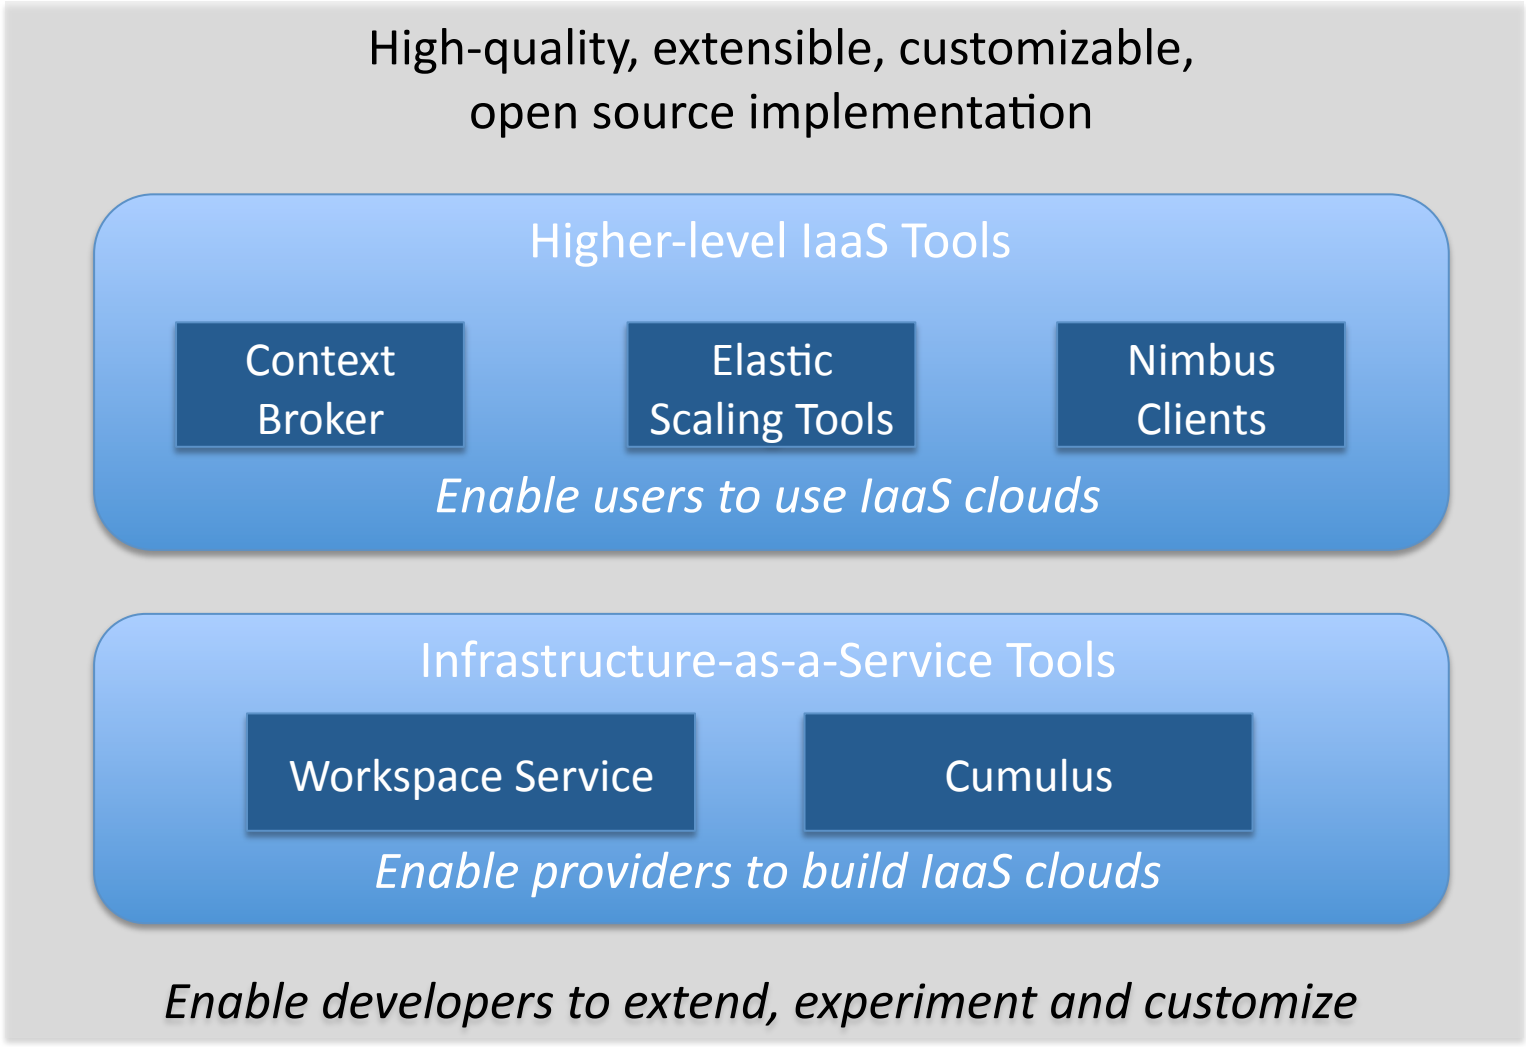

IaaS

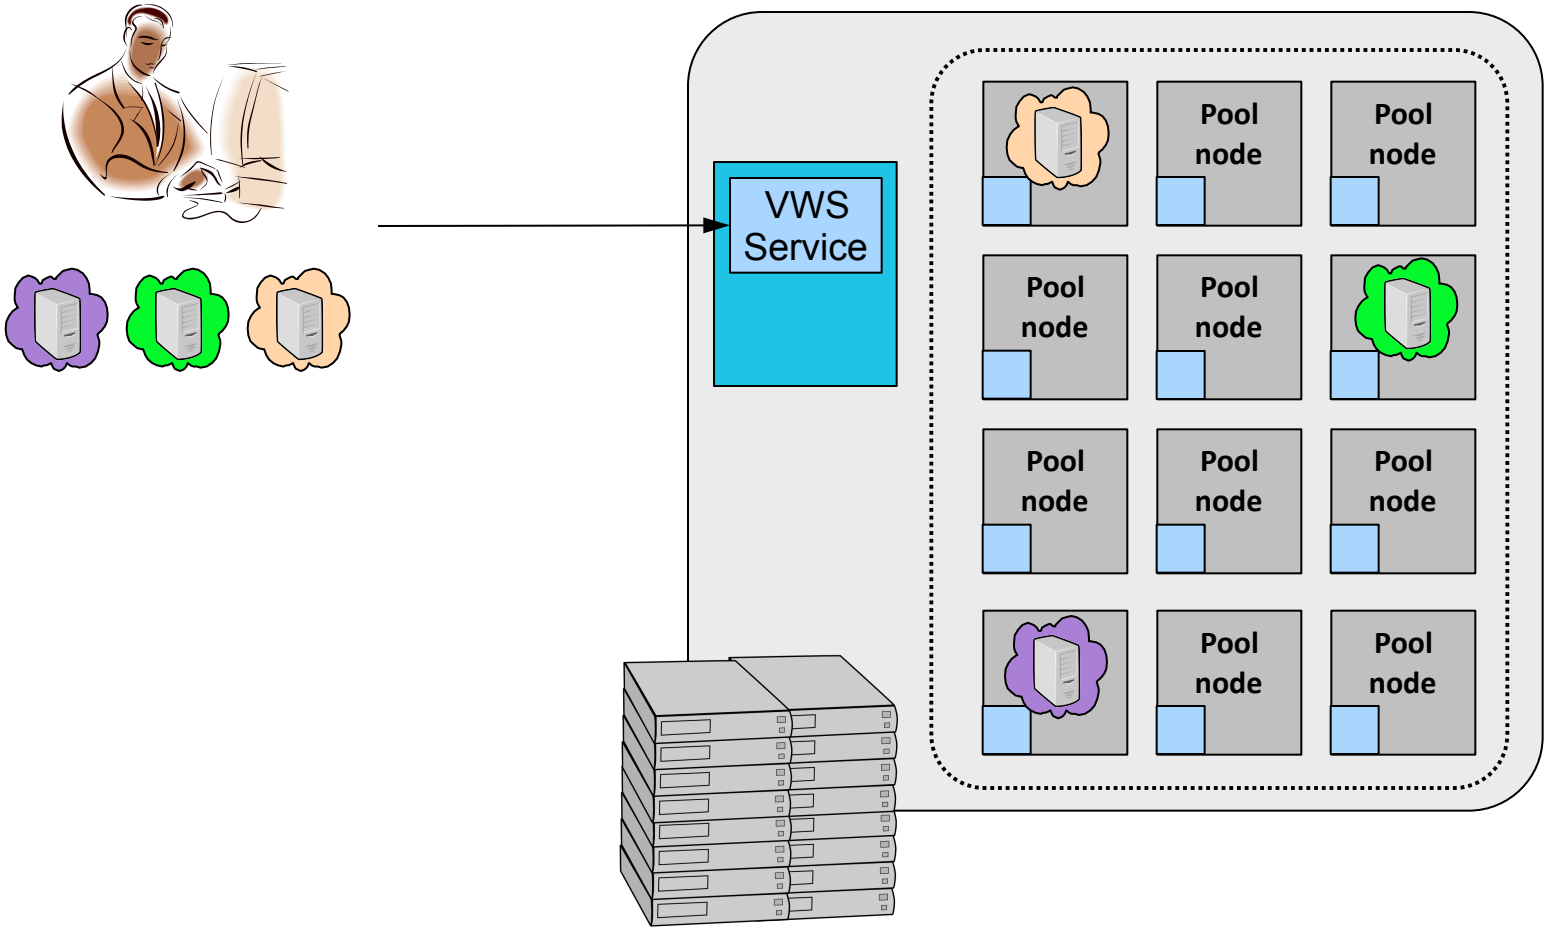
VWS Service
Pool node
Pool node
Pool node
Pool node
Pool node
Pool node
Pool node
Pool node
Pool node
Pool node
Pool node

IaaS

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.

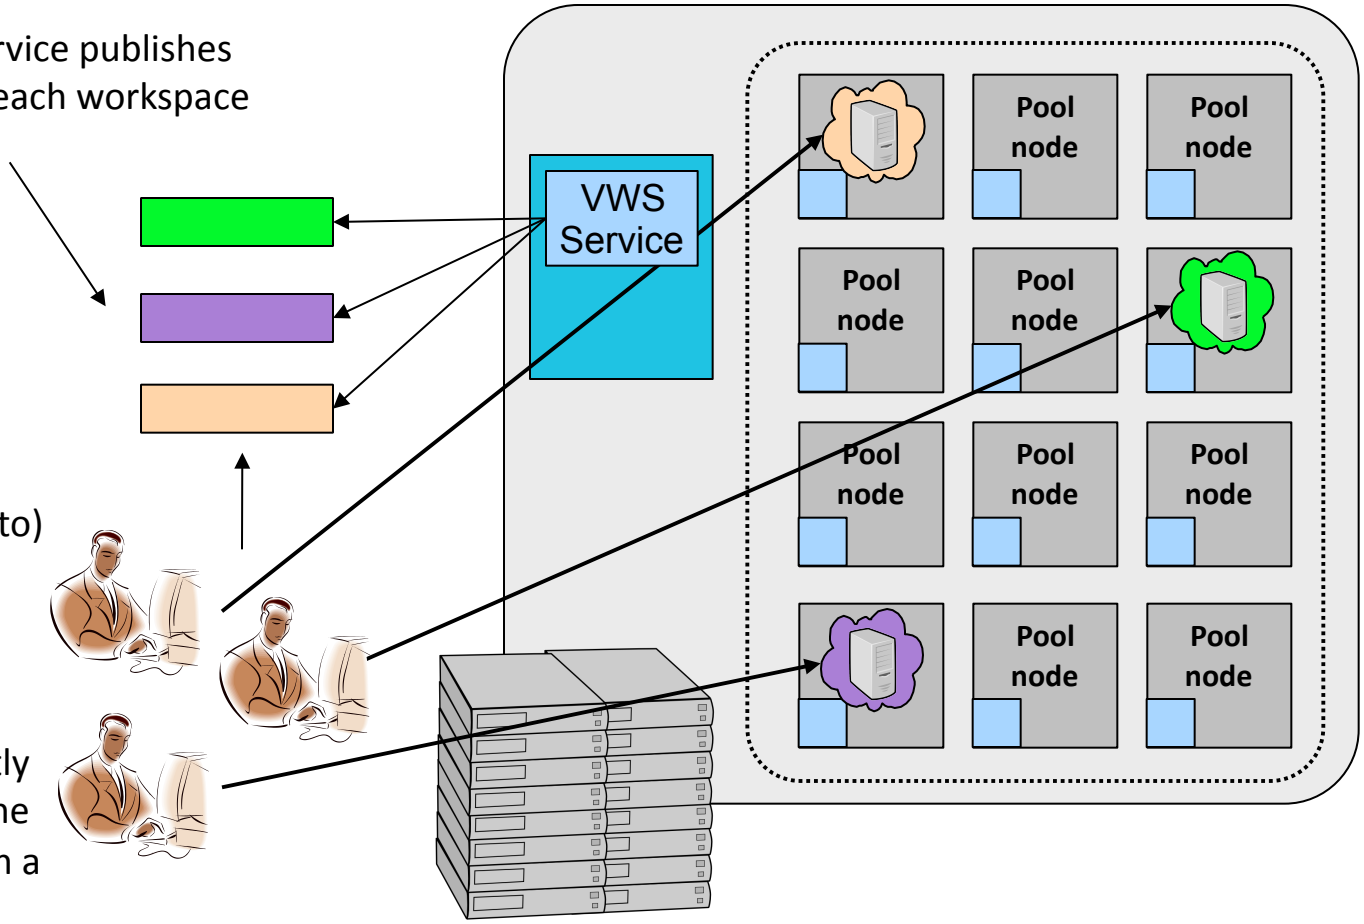

Nimbus IaaS

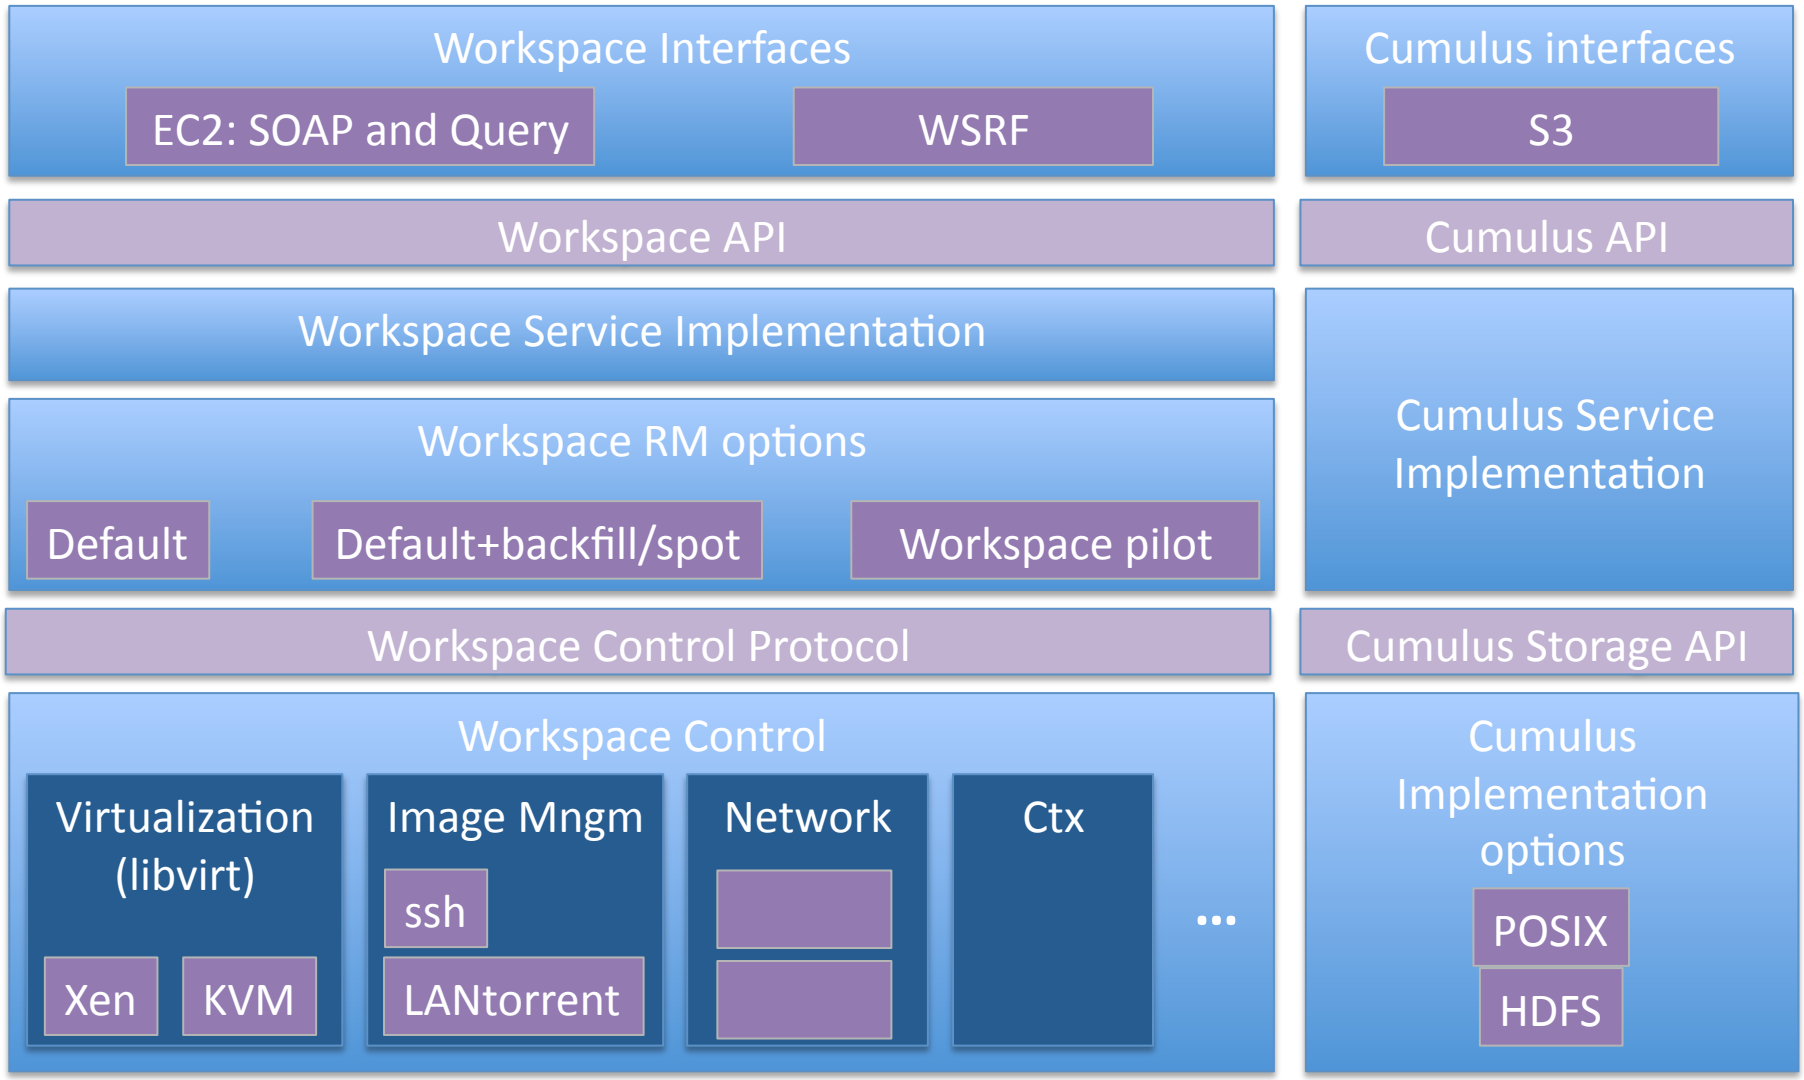
Workspace Interfaces
EC2: SOAP and Query
WSRF
Cumulus interfaces
S3
Workspace API
Cumulus API
Workspace Service Implementation
Cumulus Service Implementation
Workspace RM options
Default
Default+backfill/spot
Workspace pilot
Workspace Control Protocol
Cumulus Storage API
Workspace Control
Virtualization (libvirt)
Xen
KVM
Image Mngm
ssh
LANtorrent
Network
Ctx
...
Cumulus Implementation options
POSIX
HDFS

Working with Hybrid Clouds

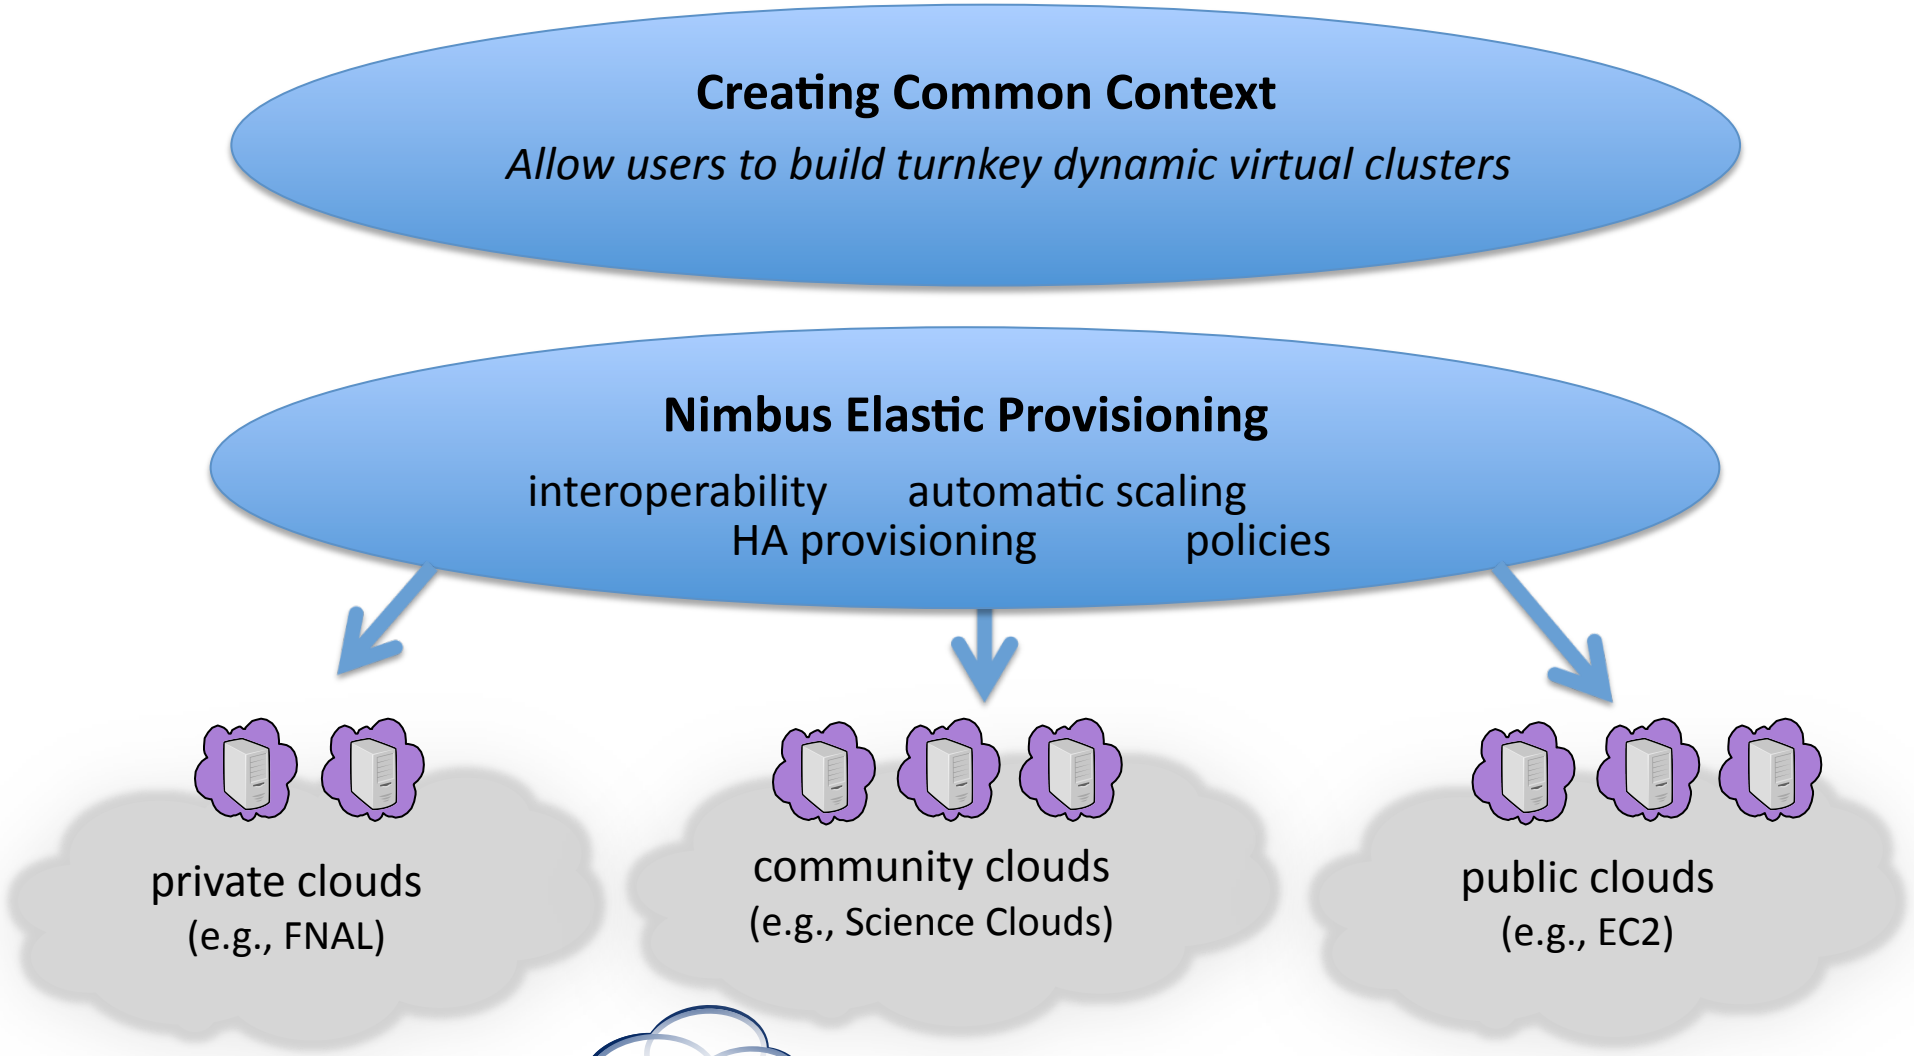

Cloud Computing for Science: A Personal Perspective

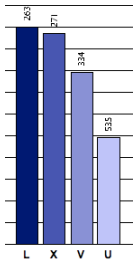

Xen released

EC2 released

First STAR production run on EC2

Science Clouds available

OOI starts

Experimental Clouds for Science

2004

2006

2008

2010

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

First Nimbus release

Nimbus Context Broker release

Applications and Ecosystem:

- Nimbus IaaS deployment
- Nimbus higher-level tools
- Ecosystem tools

Scientific Cloud Resources

- Science Clouds
  - UC, UFL, Wispy@Purdue
  - ~300 cores
- Magellan
  - DOE cloud @ ANL&LBNL
  - ~4000 cores@ANL
- FutureGrid
  - ~6000 cores
- DIAG =
  - Data Intensive Academic Grid
  - U of Maryland School of Medicine in Baltimore
  - ~1200-1500 cores
- Outside of US:
  - WestGrid, Grid5000

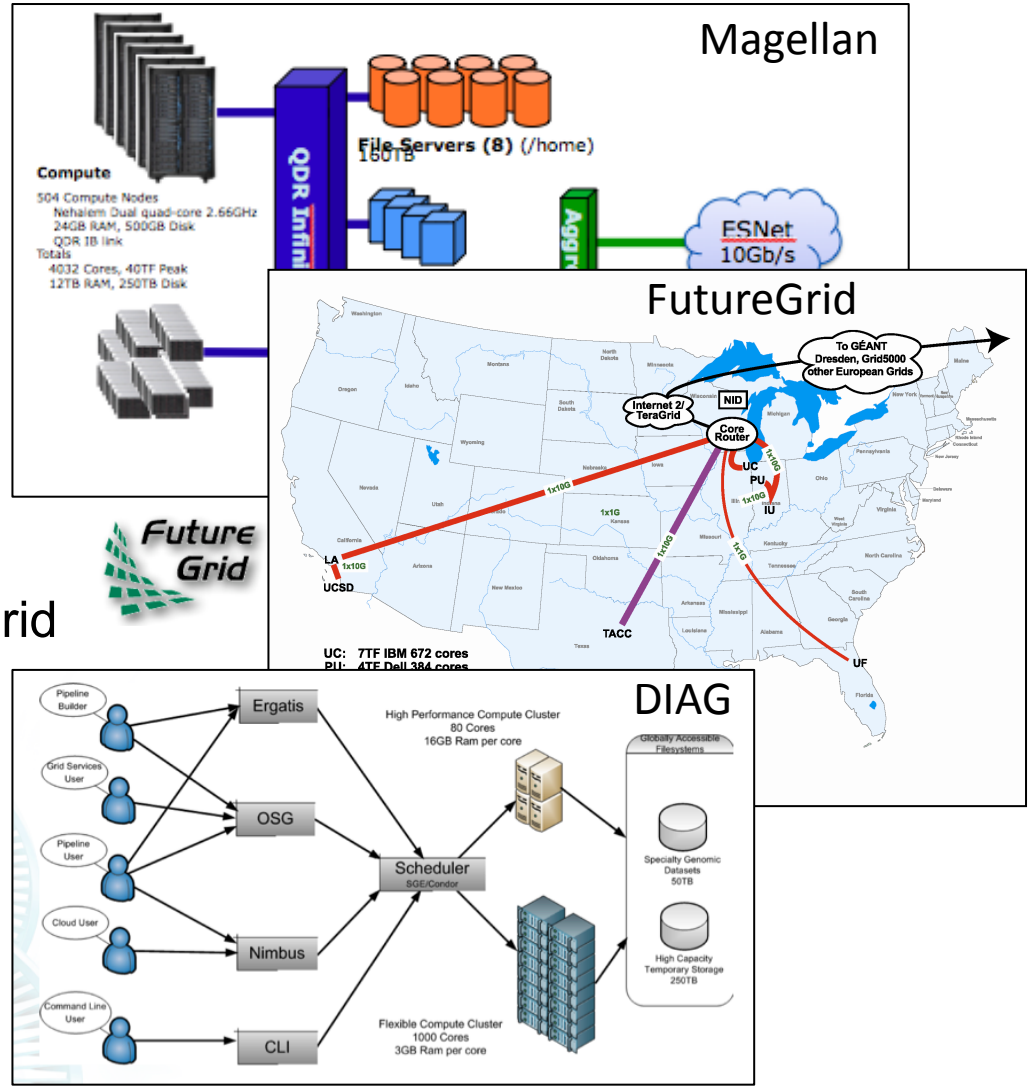

Work by Jerome Lauret (BNL) et al.

- STAR: a nuclear physics experiment at Brookhaven National Laboratory
- Virtual OSG clusters
- Science Clouds -> EC2 runs
- Production runs on EC2 since 2007
- The Quark Matter 2009 deadline: producing just-in time results
- The issue of cost

Priceless?

- Compute costs: $ 5,630.30
  - Fdsf 300+ nodes over ~10 days,
  - Instances, 32-bit, 1.7 GB memory:
    - EC2 default: 1 EC2 CPU unit
    - High-CPU Medium Instances: 5 EC2 CPU units (2 cores)
  - ~36,000 compute hours total
- Data transfer costs: $ 136.38
  - Small I/O needs : moved <1TB of data over duration

s: $ 4.69

- The people: STAR @ MIT – *Adam Kocoloski, Jan Balewski*, Mathew Walker
- Our test
  - A 100 jobs, week long simulation cost ~ $1,510
  - A year long CPU @ 100 jobs saturation ~ 79k$
  - EC2+Nimbus
    - 300+ nodes for 10 days in 2008 (non-optimized) ~ $5,600

OCEAN OBSERVATORIES INITIATIVE

- Large NSF-funded observatory with requirements for adaptive, reliable, elastic computing
- Strategy:
  - Private Nimbus regional clouds -> commer[illegible]uds
  - Highly Availab[illegible]ices that provisi[illegible]on many cl[illegible]eed
  - Signif[illegible] infra[illegible] and se[illegible]ent based on th[illegible]
- Statu[illegible]
  - Scalability and reliability tests on 100s of EC2, FutureGrid and Magellan resources
  - HA elastic services release in Spring 2011

Trail-blazing project

Sky Computing

Work by Pierre Riteau et al, University of Rennes 1

- Sky Computing = a Federation of Clouds
- Strategy:
  - Combine resources obtained in multiple Nimbus clouds in FutureGrid and Grid' 5000
  - Deployed a virtual cluster of over 1000 cores on Grid5000 and FutureGrid – largest ever of this type
  - Combine Context Broker, ViNe, fast image deployment
- Grid'5000 Large Scale Deployment Challenge award
- Demonstrated at OGF 29 06/10
- TeraGrid '10 poster

"Sky Computing"
IEEE Internet Computing, September 2009

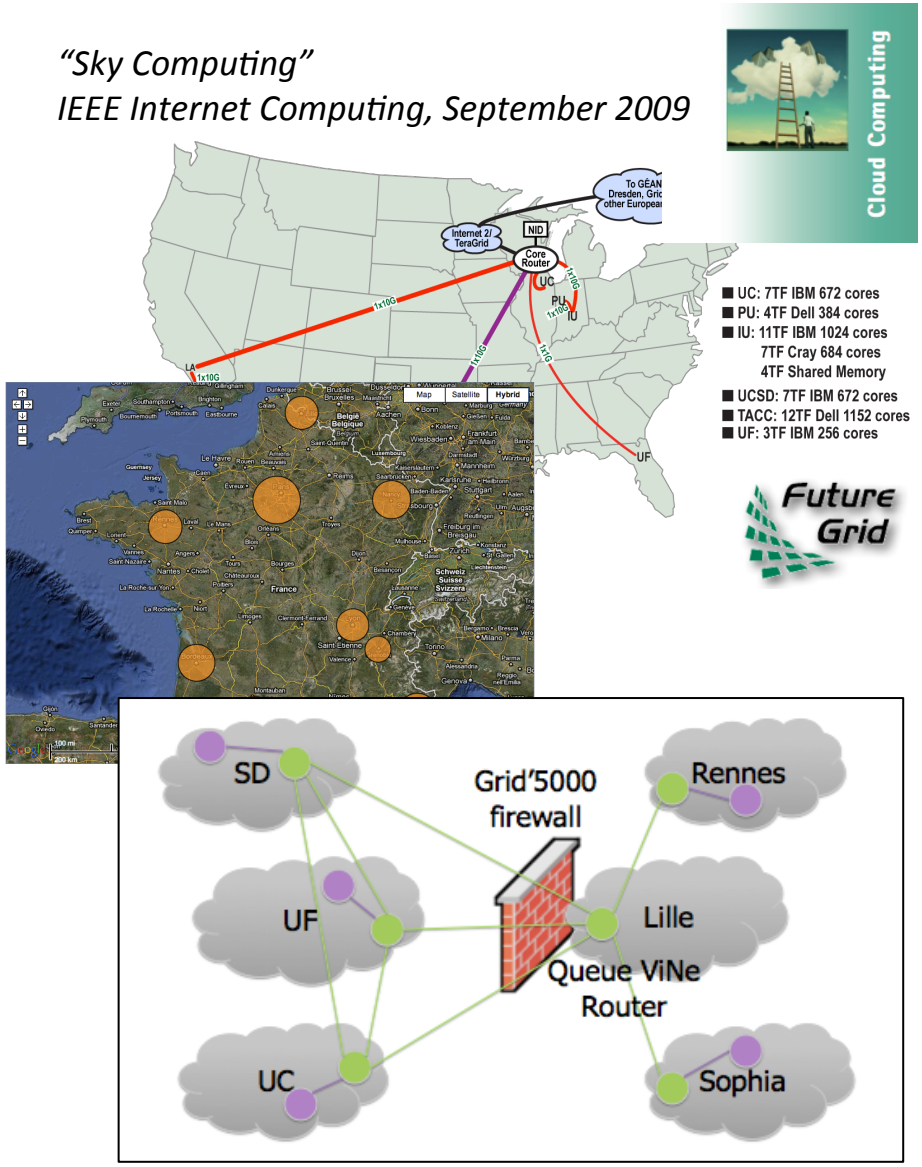

BaBar

Canadian Efforts

University of Victoria NRC·CNRC

Work by the UVIC team

- BarBar Experiment at SLAC in Stanford, CA
- Using clouds to simulating electron-positron collisions in their detector
- Exploring virtualization as a vehicle for data preservation
- Strategy:
  - Distributed Nimbus clouds
  - Appliance preparation and management
  - Cloud Scheduler
- Running production BaBar workloads

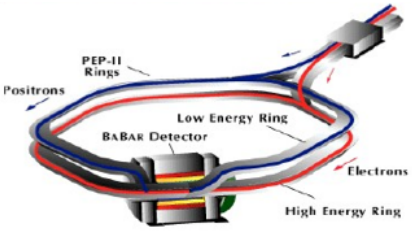

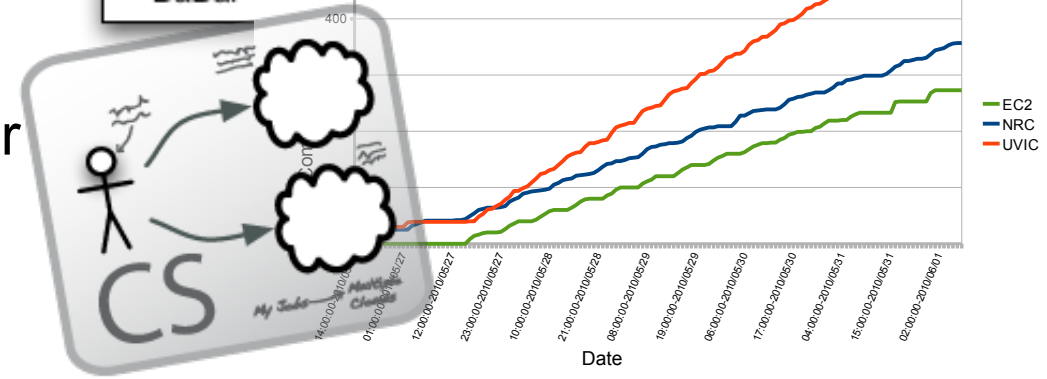

CANFAR
Canadian Advanced Network for Astronomical Research

University of Victoria NRC·CNRC

Work by the UVIC team

- Provide infrastructure for six observational astronomy survey projects
- Strategy:
  - Running on a Nimbus cloud on WestGrid
  - Dynamic Condor pool for astronomy
  - Appliance creation and management
- Status:
  - MACHO experiment Dark Matter search
  - In production operation since July 2010

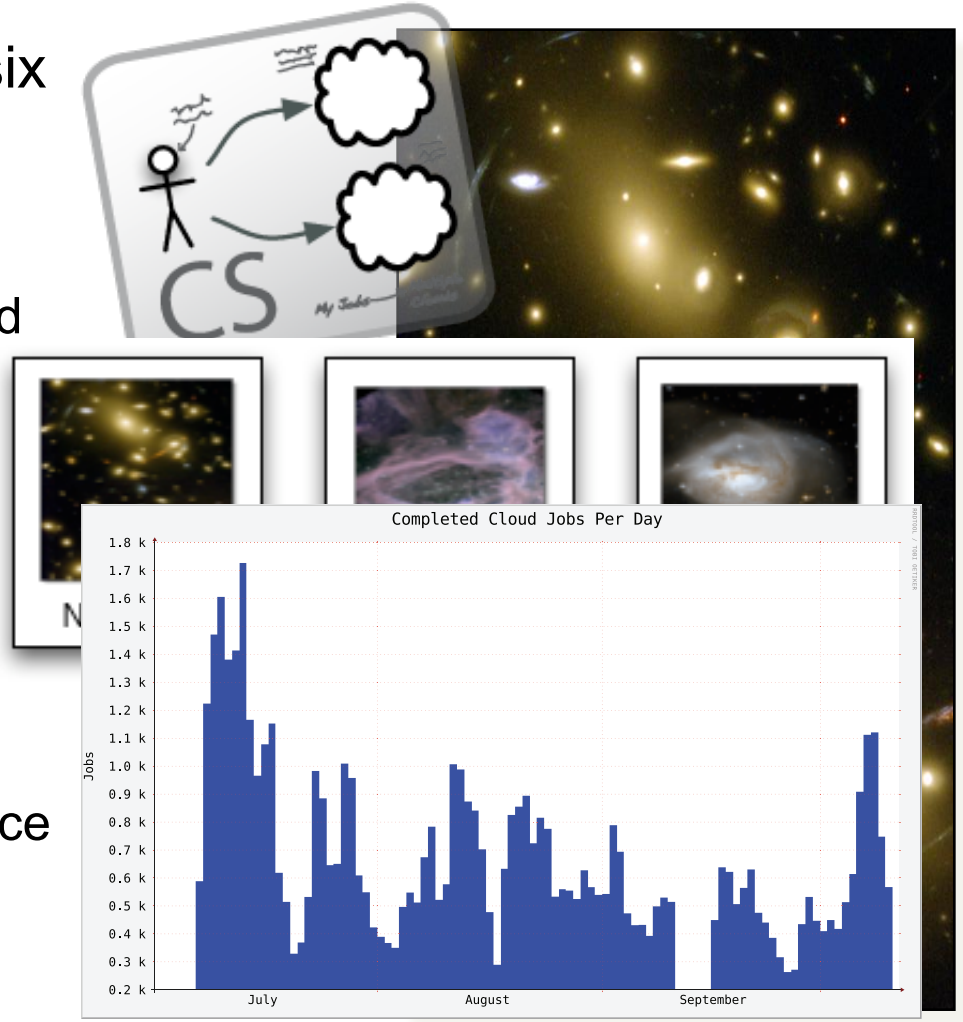

CloVR

Sam Angiuoli
Institute for Genome Sciences
University of Maryland School of Medicine

- The emergent need for processing
- A virtual appliance for automated and portable sequence analysis
- Strategy:
  - Running on Nimbus Science Clouds, Magellan and EC2
  - A platform for building appliances representing push-button pipelines
- Impact
  - From desktop to cloud
  - http://clovr.org

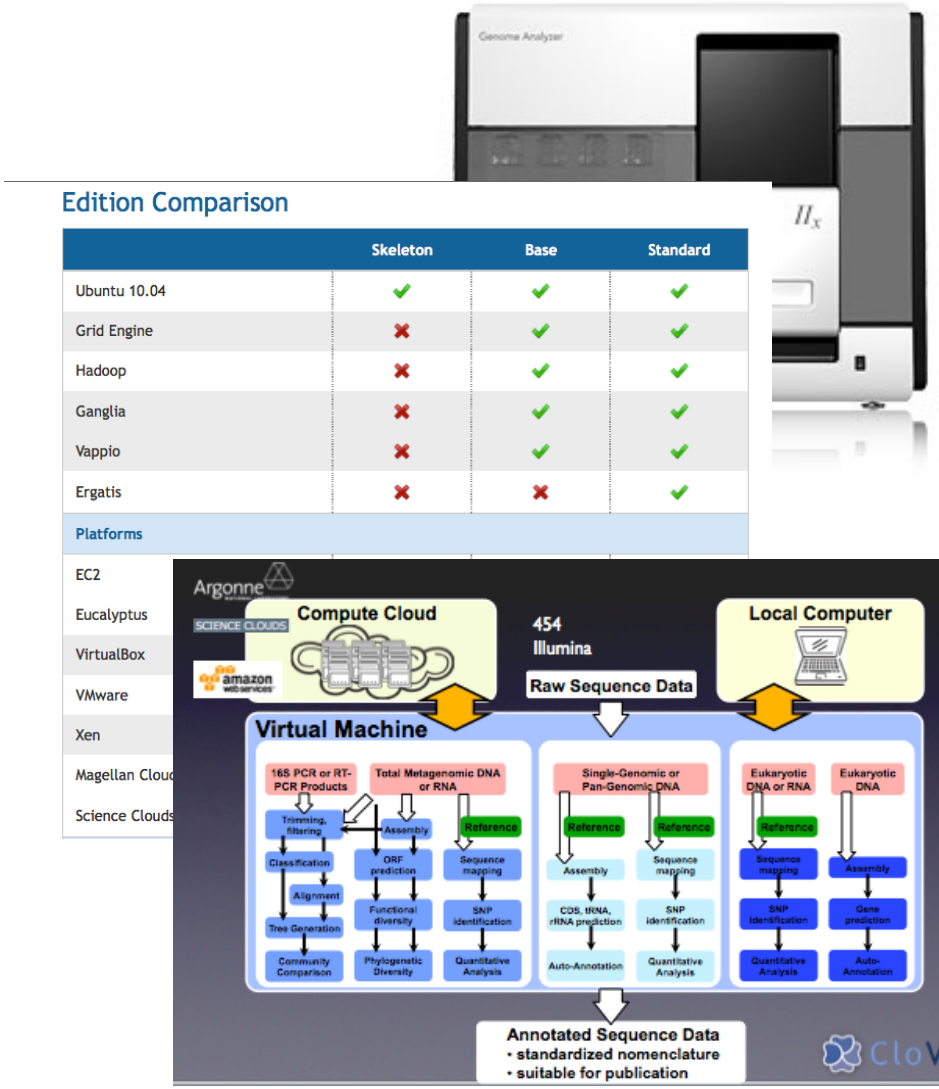

Blessings

- Deploy custom, user-owned and user-controlled environments on remote resources
- On-demand access
- Elastic processing
- Growth and cost management
- Capital expense -> operational expense

Challenges (Some of)

- Appliance management
- Lack of reliability
- Elasticity, but how?
- Performance of deployment
- Clouds are NOT high performance!
- High cost
- Lack of markets

*Disclaimer: the security challenge has been left to Alvaro to solve!

Appliance Management

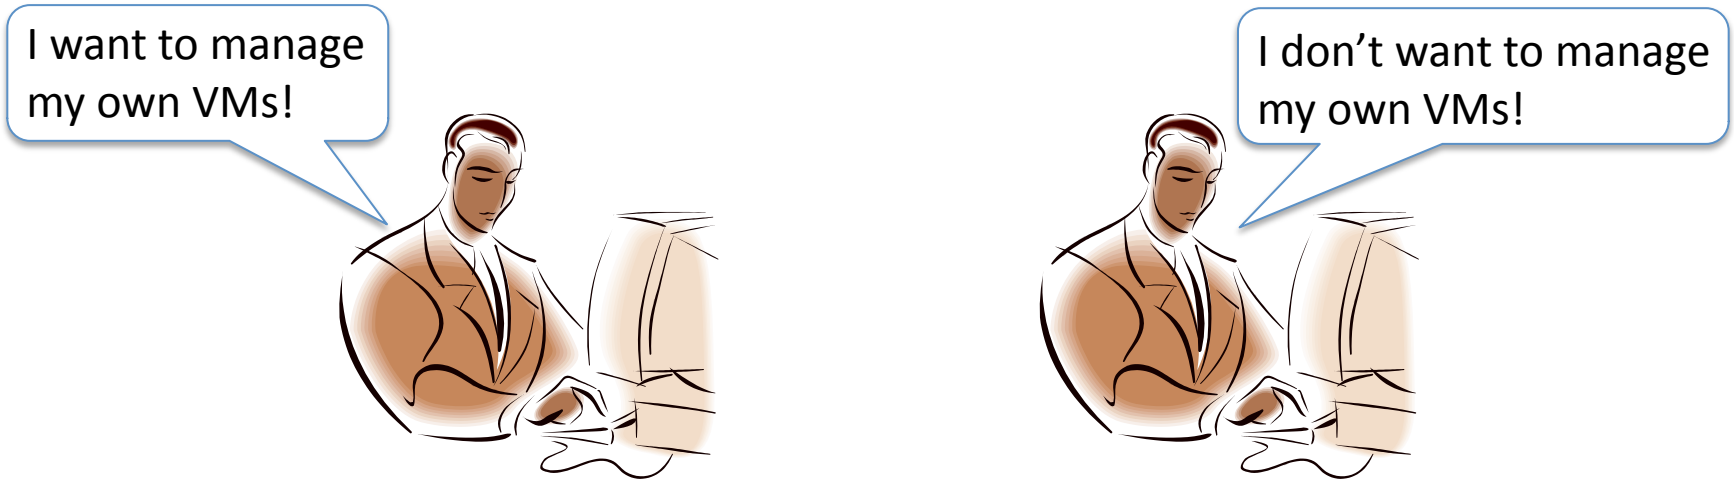

- Many users are not skilled in appliance development yet require the control
- The emergence of community appliance management and maintenance
  - Clovr, BarBar and Canfar @ UVIC, CernVM, and others
- A new role/job description
- Forward looking applications: data preservation

Elasticity, Reliability and Failure

Elasticity and reliability are different sides of the same coin.

- 2008: The ALICE proof-of-concept
- 2009: ElasticSite prototype
- 2009: OOI pilot

Need for generic, HA, elastic service model

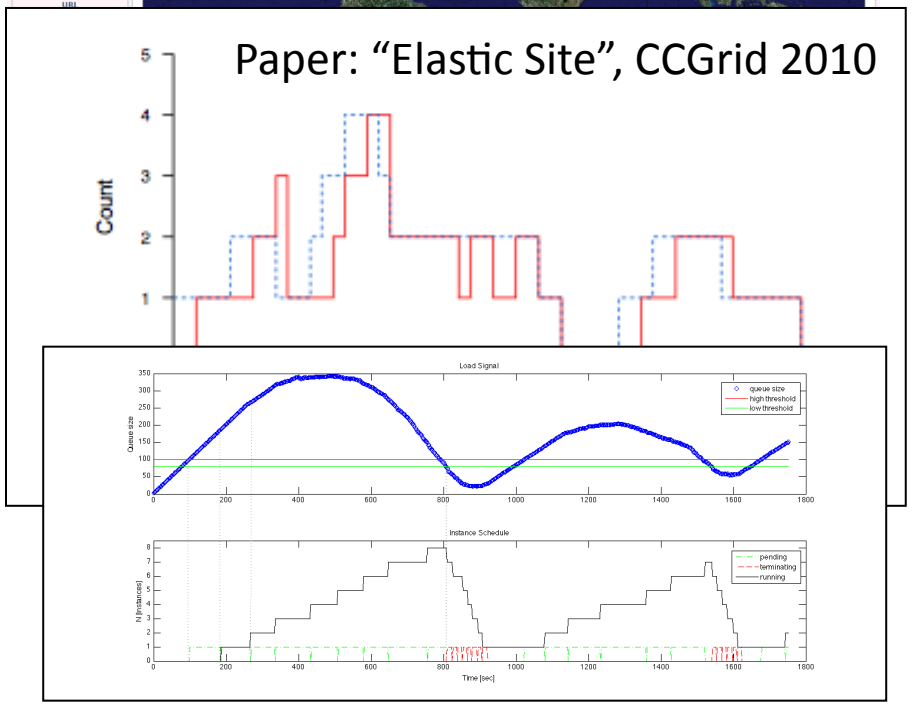

Elasticity, Reliability and Failure

- Assumption: a workload queue
  - ALiEn, PBS, AMQP,…
- React to sensor information
  - Queue properties a sensor
- Scale to demand
  - Across different cloud providers
  - Use contextualization to integrate machines across hybrid clouds
  - Highly Available
  - Scalable: latest tests scale to 100s of nodes on EC2, target is thousands
- Release in early 2011
  - Customizable to input, policy, provider, etc.

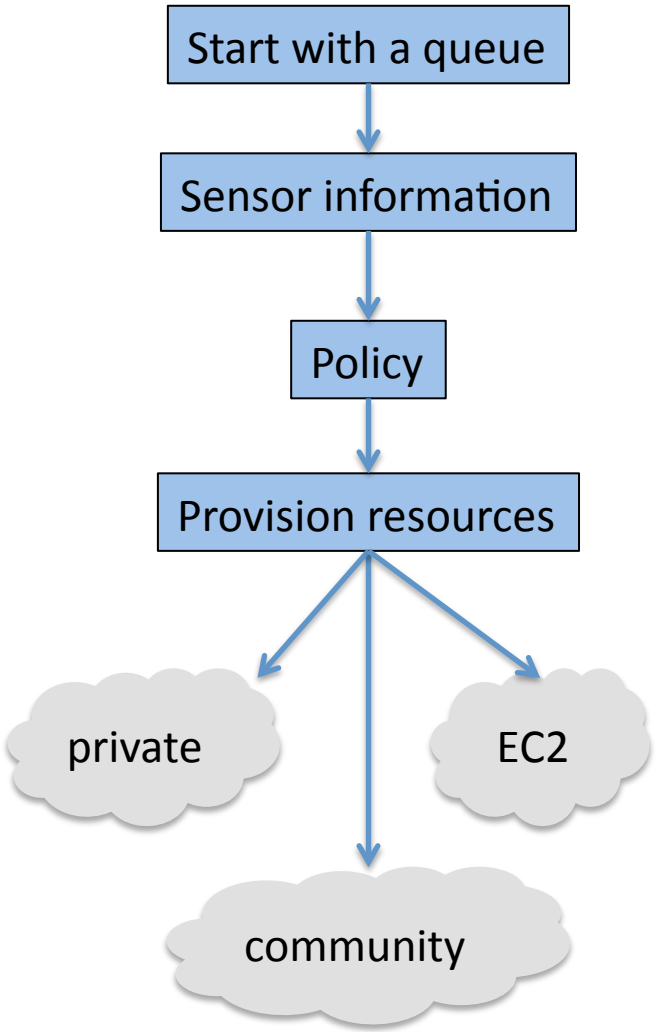

Deployment Performance

- Moving images is the main component of VM deployment
- LANTorrent: the BitTorrent principle on a LAN
- Streaming
- Minimizes congestion at the switch
- Detecting and eliminating duplicate transfers
- Benefit: a thousand VMs in 10 minutes
- Nimbus release 2.6

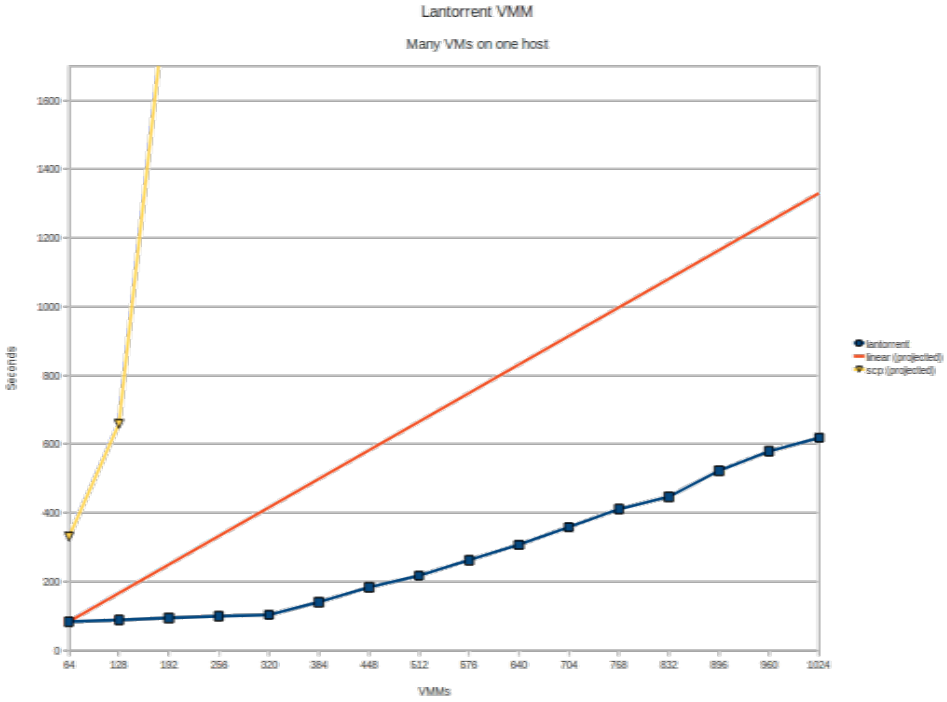

Preliminary data using the Magellan resource
At Argonne National Laboratory

Execution Performance

- Virtualization overhead: myth or reality?
- Ingredients of performance
  - Hardware
  - Configuration
  - Stability
- Things are looking up!
- From performance…
- …to price-performance
  - Instances for science
- Availability @ price point

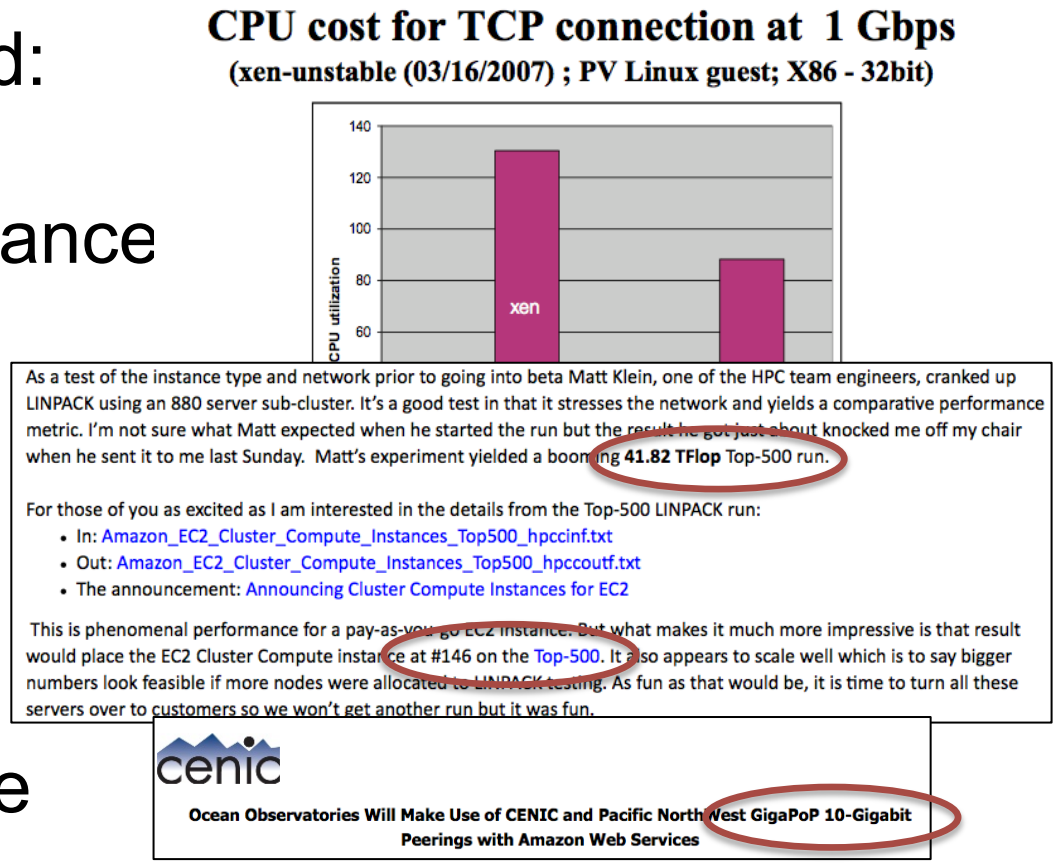

Cost, Utilization, and Price

- Most science today is done in batch
  - Not very responsive…
  - … but very efficient!
- On-demand catch-22 for providers:
  - You can overprovision (Expensive!)
  - Or you can reject requests (Not really on-demand)
- Utilization -> Cost -> Price

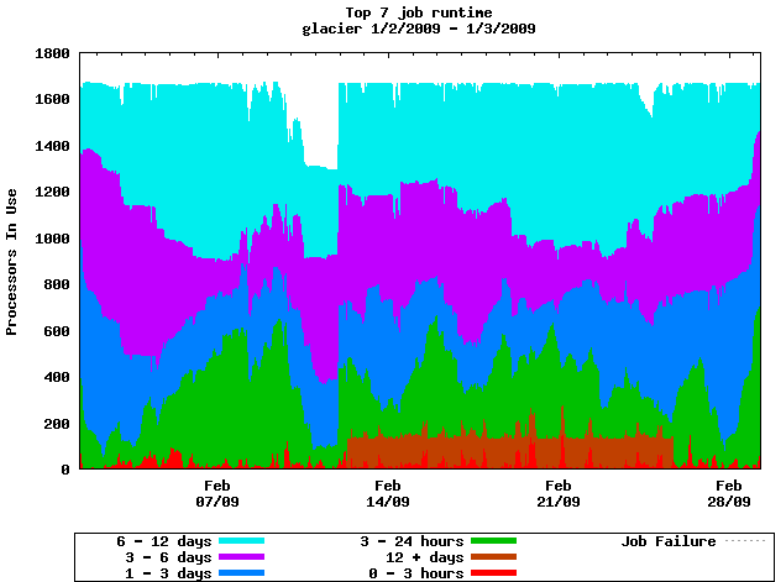

courtesy of Rob Simmonds, example of WestGrid utilization

Cost, Utilization, and Price

- Solution 1:
  - Backfill with volunteer VMs
- Benefits:
  - Up to 100% utilization!
  - Significantly lower cost
- Solution 2:
  - Spot pricing
  - Support for auctions
- Open Source community contribution
- Preparing for running of production workloads on FG @ U Chicago
- Extension to Nimbus Workspace Service RM to be available in Nimbus release 2.7

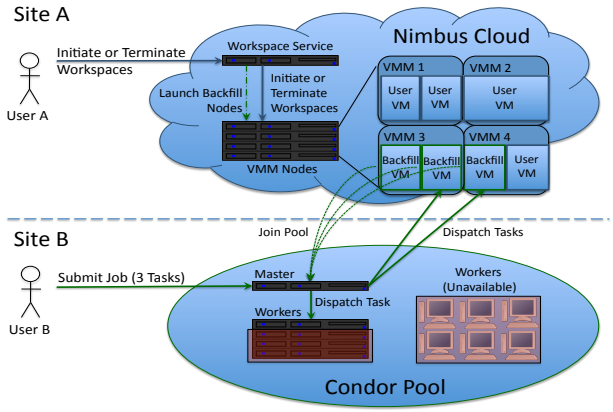

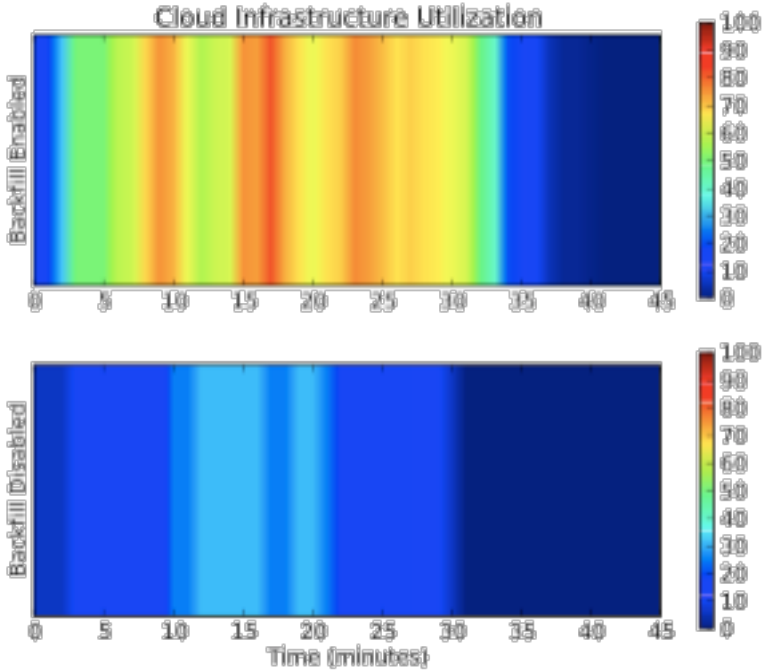

Cloud Markets (Lack Thereof)

- Are cloud resources fungible?
- IaaS API standards:
  - Cloud standards: OCCI (OGF), OVF (DMTF), and many more…
  - www.cloud-standards.org
- IaaS API adapters
  - Deltacloud, jcloud, libcloud, and many more…
- VM image standards….
- …and adapters
  - rBuilder, BCFG2, CohesiveFT, Puppet, and many more…
- Comparing value….
  - Comprehensive
  - Current
  - And public!

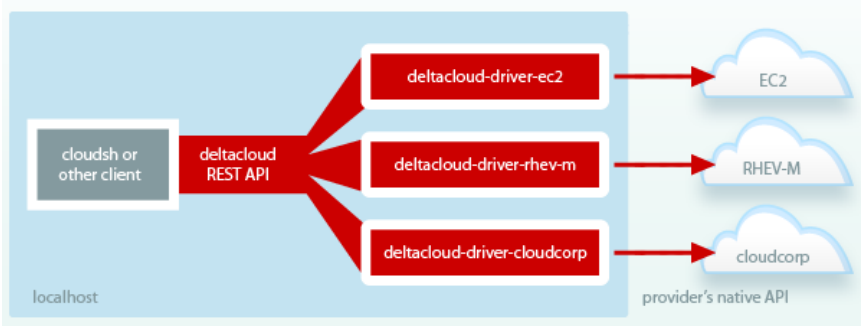

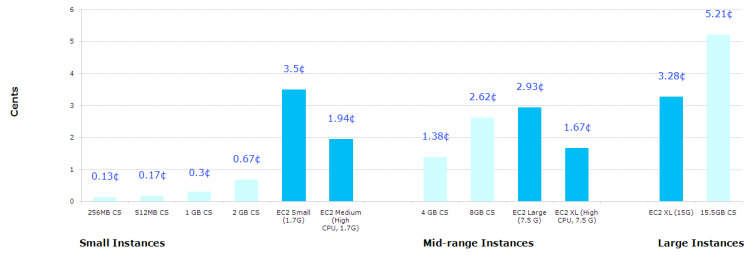

The Bitsource: CloudServers vs EC2 LKC Cost by Instance

Nimbus Collaboration

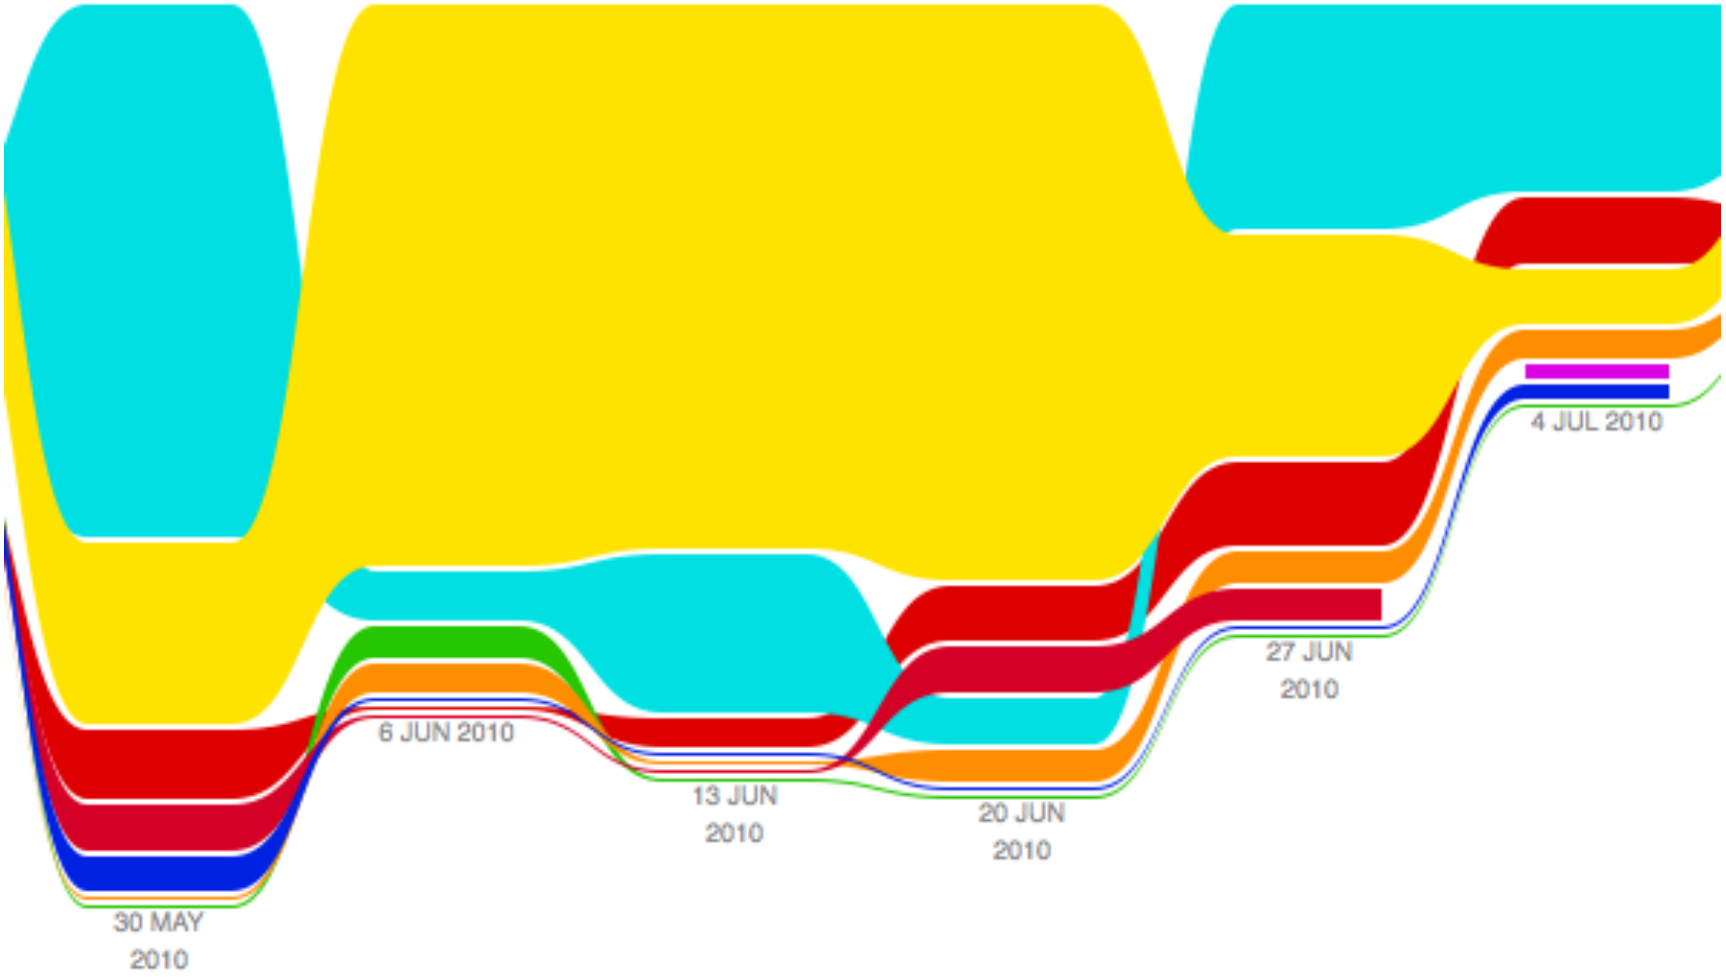

The Nimbus Team

- Project lead: Kate Keahey, ANL&UC
- Comitters:
  - Tim Freeman - University of Chicago
  - Ian Gable - University of Victoria
  - David LaBissoniere - University of Chicago
  - John Bresnahan - Argonne National Laboratory
  - Patrick Armstrong - University of Victoria
  - Pierre Riteau - University of Rennes 1, IRISA
- Github Contributors:
  - *Tim Freeman, David LaBissoniere, John Bresnahan, Pierre Riteau, Alex Clemesha, Paulo Gomez, Patrick Armstrong, Matt Vliet, Ian Gable, Paul Marshall, Adam Bishop*
- *And many others*
  - *See http://www.nimbusproject.org/about/people/*

Open Source Project

- Necessary conditions
  - Open source license
  - Open code
- Sufficient conditions
  - Design for extensibility
  - Acceptance test framework
  - Social coding mechanisms
  - Inclusive team dynamic

Parting Thoughts

- Cloud computing is here to stay
- We showed that it is important to a set of applications – how can we make that set larger?
- A change of paradigm -> a change of pattern
  - New exciting modes of usage
  - New exciting opportunities and needs
- Open source rocks!