

Open Cirrus Summit 2011

Oct 12-13, Atlanta, Georgia

Building an Outsourcing Ecosystem for Science




Kate Keahey

keahey@mcs.anl.gov

Argonne National Laboratory
Computation Institute, University of Chicago

Outsourcing Computing for Science



Control over environment

- Complexity
- Consistency

Control over availability

- Many real-time scenarios
- Naturally bursty character



NIMBUS

www.nimbusproject.org

Introducing Nimbus

High-quality, extensible, customizable,
open source implementation

Nimbus Platform

Context
Broker

Cloudinit.d

Elastic
Scaling Tools

Enable users to use IaaS clouds

Nimbus Infrastructure

Workspace
Service


Cumulus

Enable providers to build IaaS clouds

*Enable developers to extend,
experiment and customize*



On the Nimbus Team...



Applications and Patterns



Providing infrastructure outsourcing for science

Applications and Patterns


10/12/11




www.nimbusproject.org

Scientific Cloud Resources

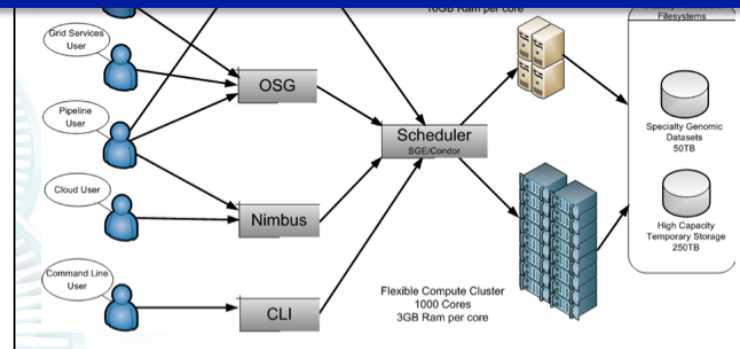
- Science Clouds
 - UC, UFL, Wispy@Purdue
 - ~300 cores
- Magellan
 - DOE cloud @ ANL & LBNL





has cloud cycles for you!

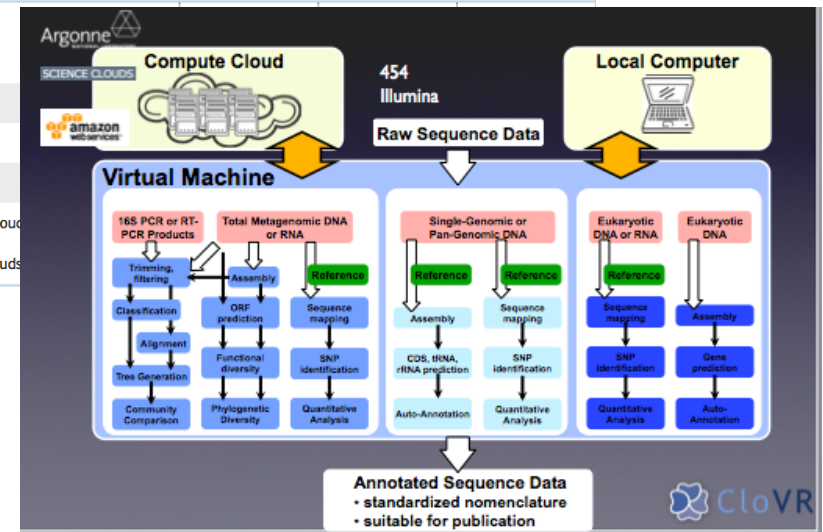
- Medicine in Baltimore
 - ~1200-1500 cores
- Outside of US:
 - WestGrid, Grid'5000






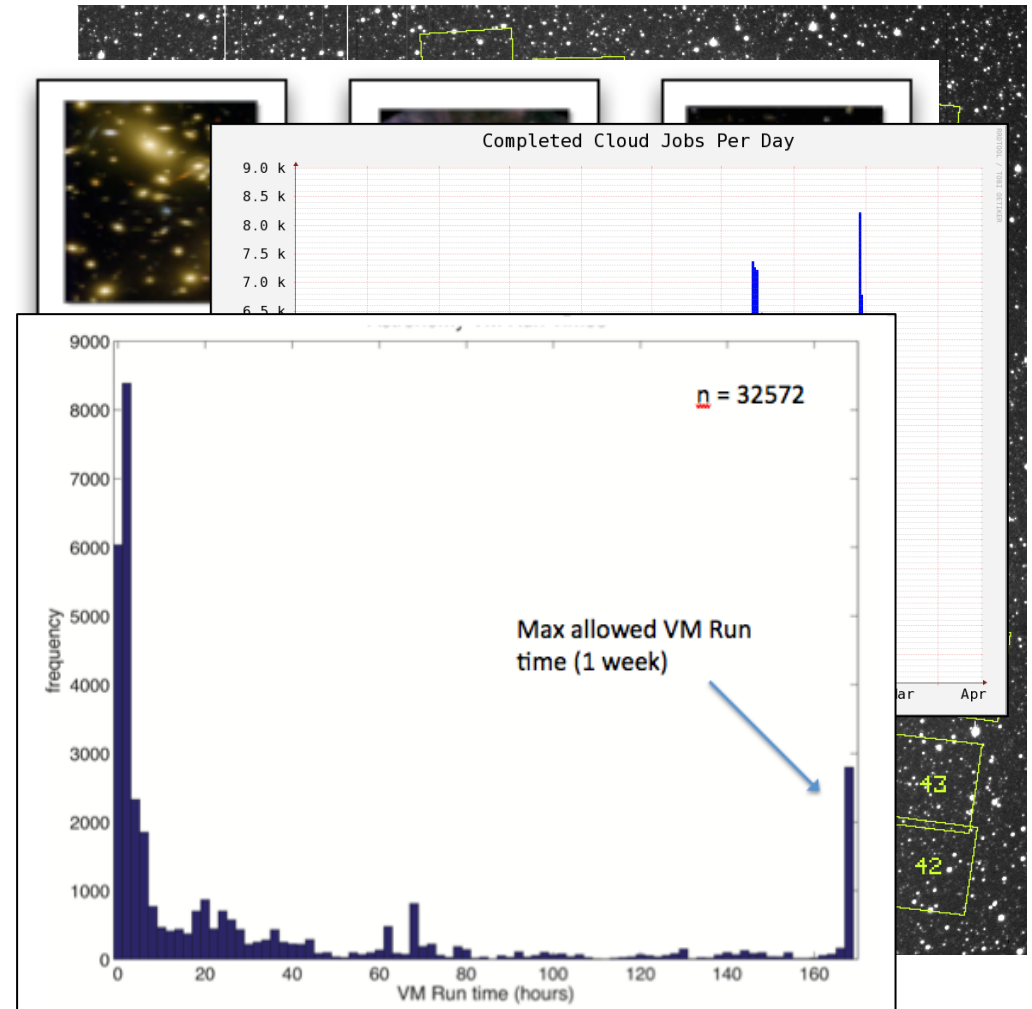
Sam Angiuoli
 Institute for Genome Sciences
 University of Maryland School of Medicine

- The emergent need for processing
- A virtual appliance for automated and portable sequence analysis
- Approach:
 - Running on user's desktop
 - Then 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>





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


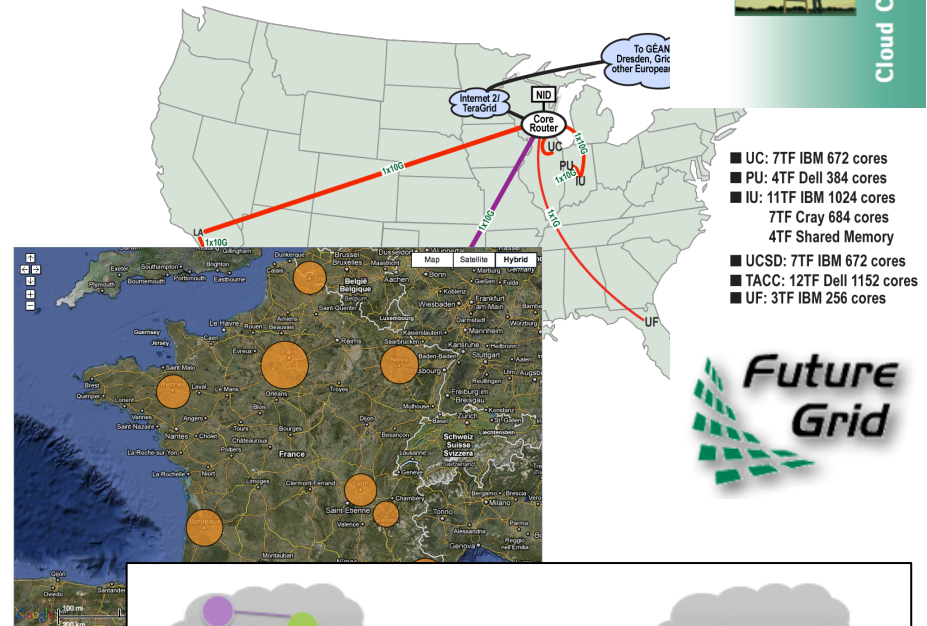


Sky Computing

Work by Pierre Riteau et al,
University of Rennes 1

- Sky Computing = a Federation of Clouds
- Approach:
 - Combine resources obtained in multiple Nimbus clouds in FutureGrid and Grid' 5000
 - Combine Context Broker, ViNe, fast image deployment
 - Deployed a virtual cluster of over 1000 cores on Grid5000 and FutureGrid – largest of this type at the time
- Grid'5000 Large Scale Deployment Challenge award
- Demonstrated at OGF 29 06/10
- TeraGrid '10 poster
- More at: www.isgtw.org/?pid=1002832

"Sky Computing"
IEEE Internet Computing, September 2009





BABAR

™ and © Neilsons, All Rights Reserved


Canadian Efforts



University of Victoria

NRC-CMRC


Work by the UVIC team



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




Resource	Cores	Notes
FutureGrid @Argonne Lab	100 Cores Allocated	Resources allocation to support BaBar
Elephant Cluster @Uvic	88 Cores	Experimental cloud cluster hosts (xrootd for cloud)
NRC Cloud in Ottawa	68 Cores	Hosts VM image repository (repoman)
Amazon EC2	Proportional to \$	Grant funding from Amazon
Hermes Cluster @Uvic	Variable (280 max)	Occasional Backfill access



STAR

Work by J. Lauret, J. Balewski and the STAR team

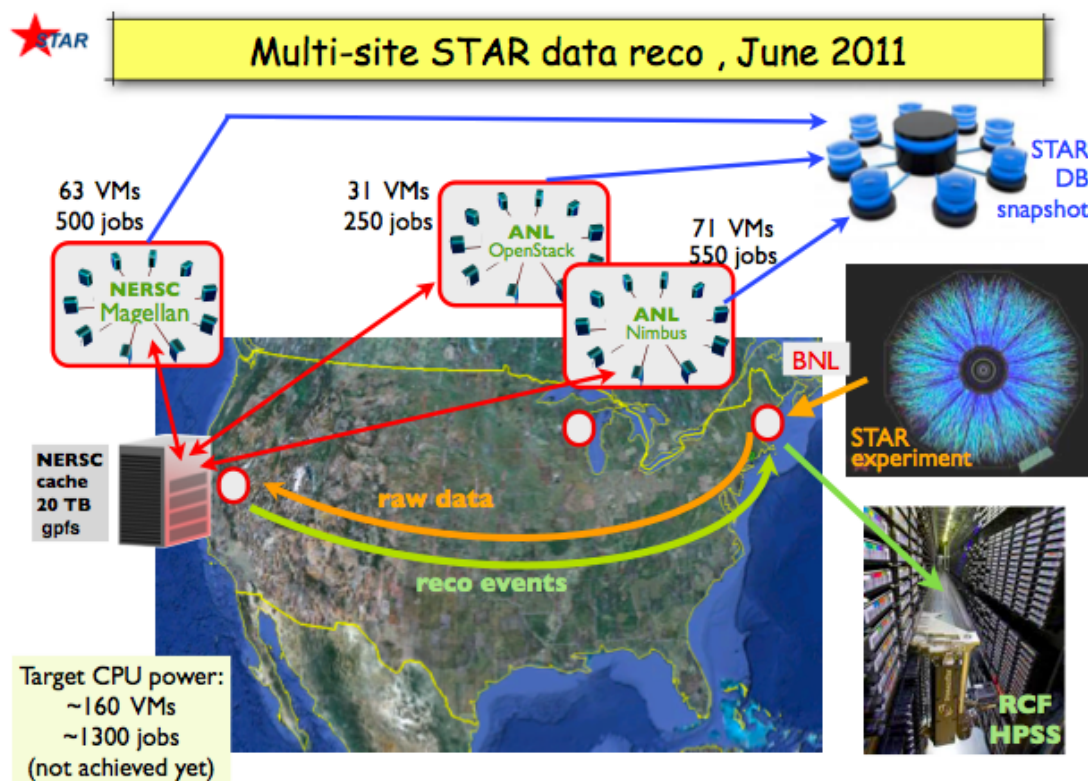



Challenge: what makes a proton spin?

Achieved timeline of W measurement in 2009



Intended timeline of W measurement in 2011



- Why run in the cloud?
 - Reduce "time to science"
 - Near real-time processing

Images courtesy of J. Balewski

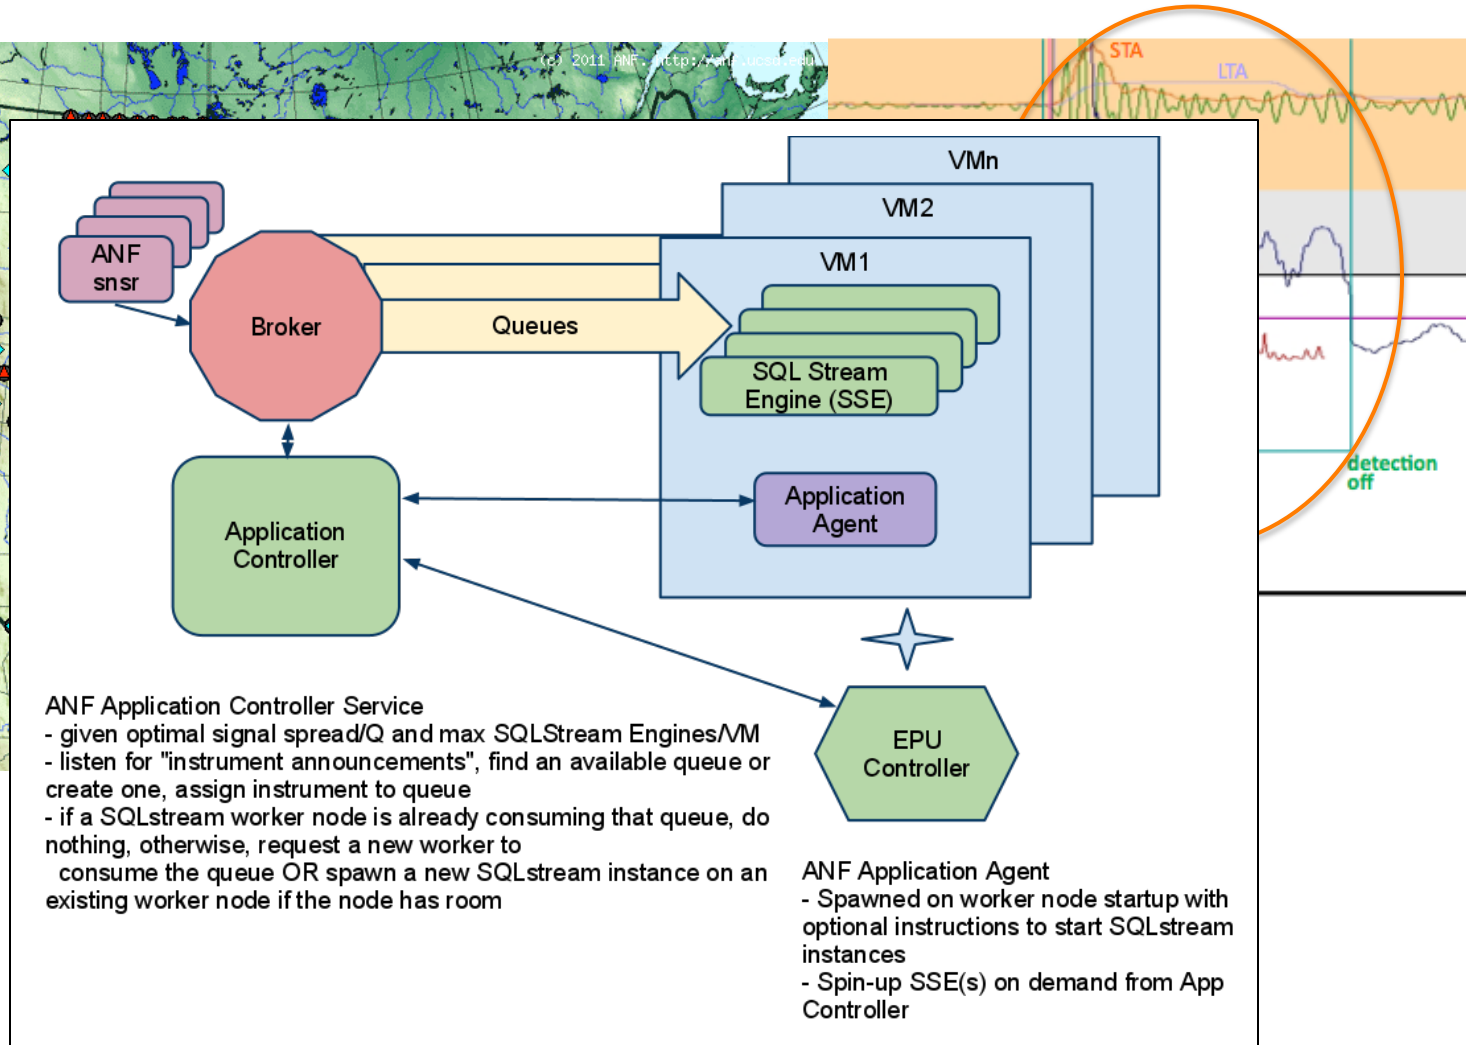
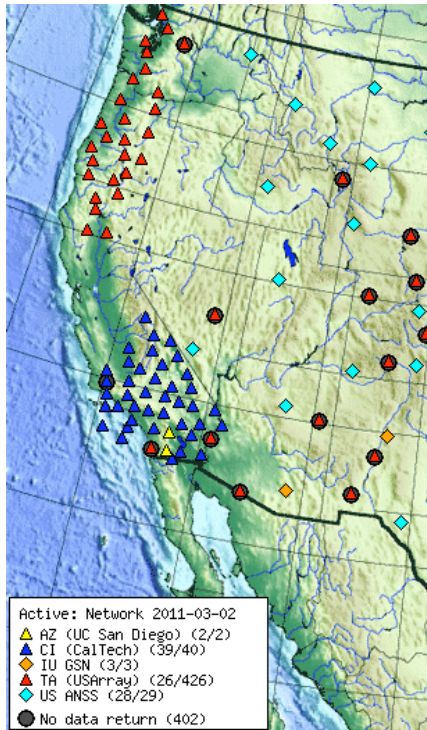


- Towards Observatory Science
- Sensor-driven processing
- Scalable and Highly Available (HA) services
- Nimbus team leading the development of Common Execution Infrastructure
- From regional Nimbus clouds to commercial clouds
- Building platform services for integrated, repeatable support for on-demand science



Array Network Facility (ANF)

Using elastic processing to operate on seismic data



Trends and Patterns

- On-demand, **elastic** processing
 - Observatories, experiments, conference deadlines, growth management...
- Multiple providers
 - Not enough cycles, but also **risk-mitigation**
- From the desktop to the cloud
 - **Escalation** pattern: seamlessly integrating more resources
- Ease of use/low barrier
 - **Automated** provisioning of infrastructure resources
- From one-offs to **production** runs
 - Steadily increasing in both size and buy-in
- Emphasis still on **loosely-coupled**
 - ... EC2 ClusterCompute virtual clusters on Top500 list

Outsourcing for Science: Building an Infrastructure Platform

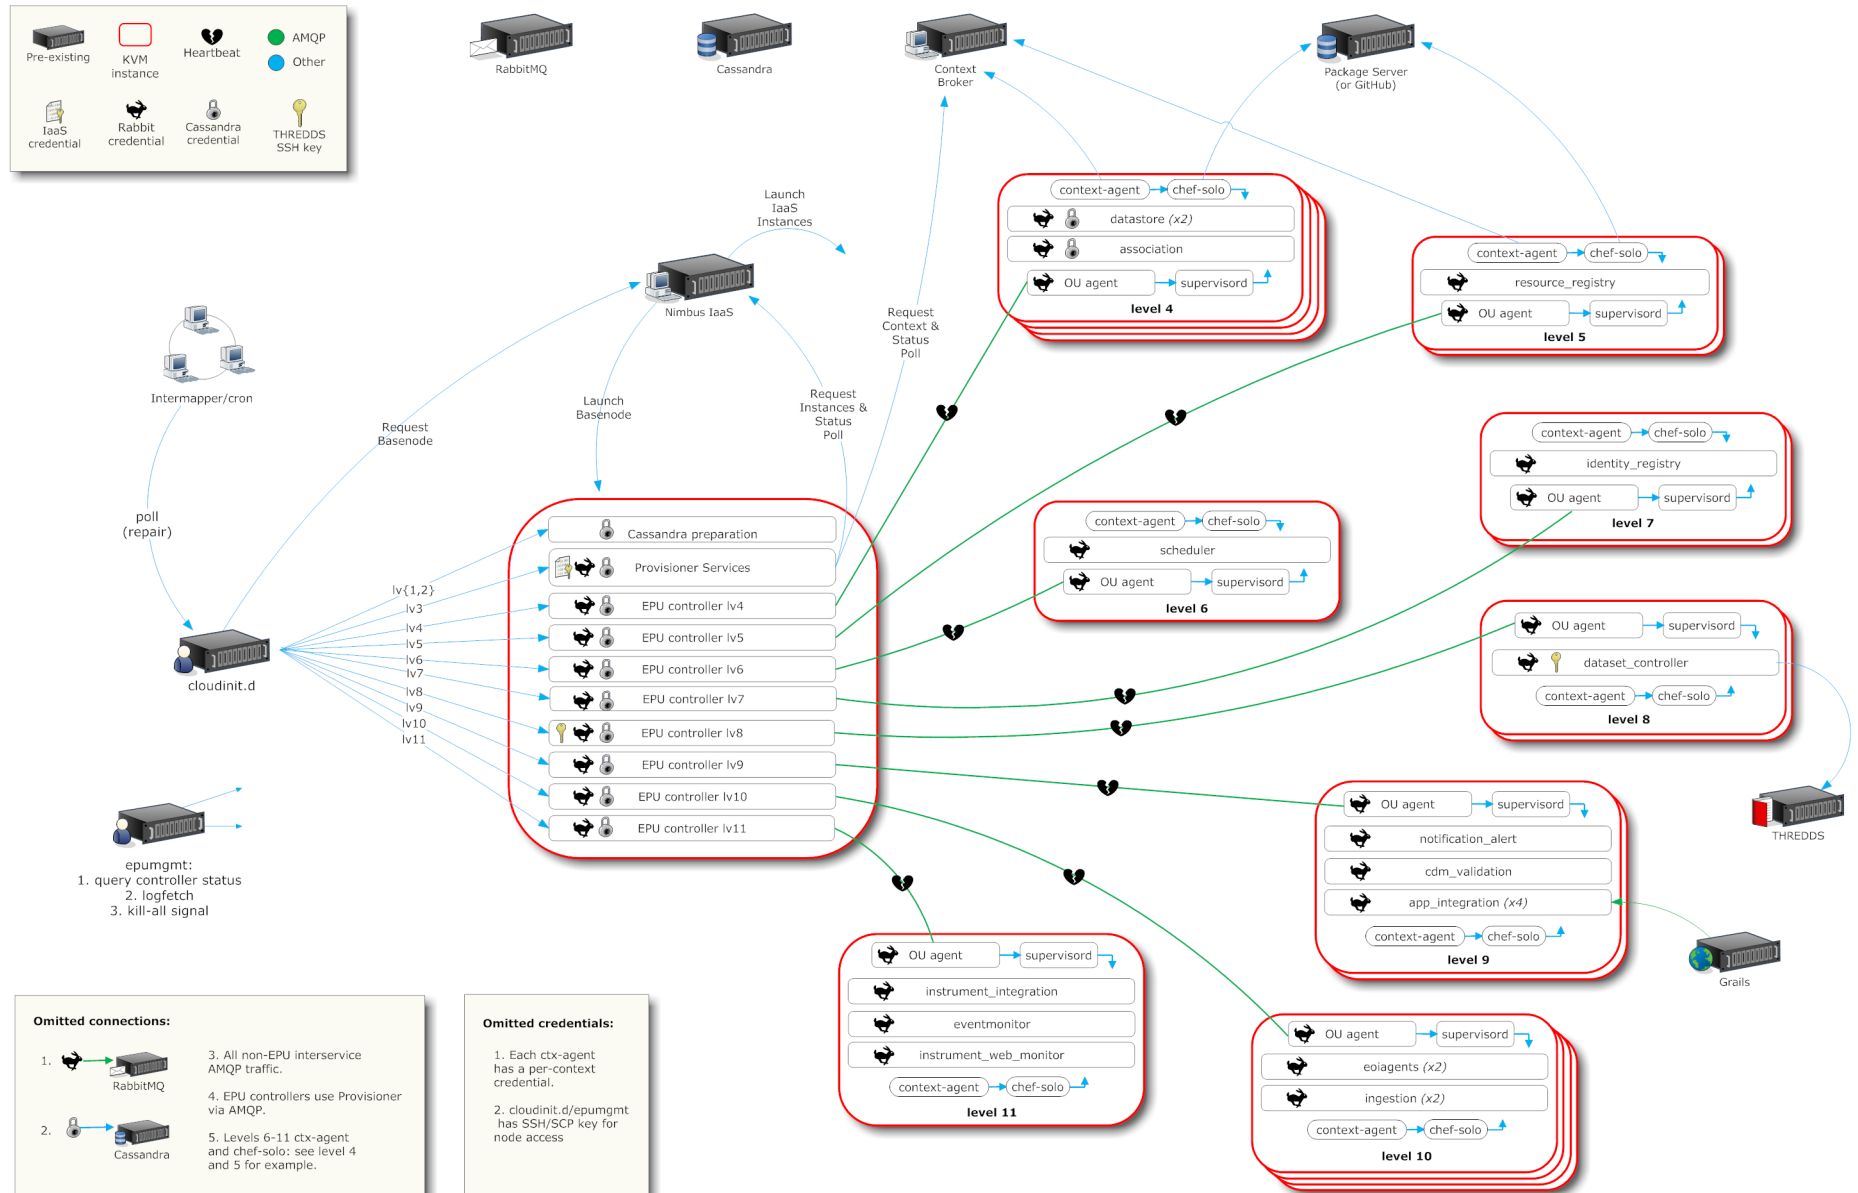
10/12/11



www.nimbusproject.org

15

A Simplified Deployment Scenario



Omitted connections:

1. [Icon] → RabbitMQ
2. [Icon] → Cassandra
3. All non-EPU interservice AMQP traffic.
4. EPU controllers use Provisioner via AMQP.
5. Levels 6-11 ctx-agent and chef-solo: see level 4 and 5 for example.

Omitted credentials:

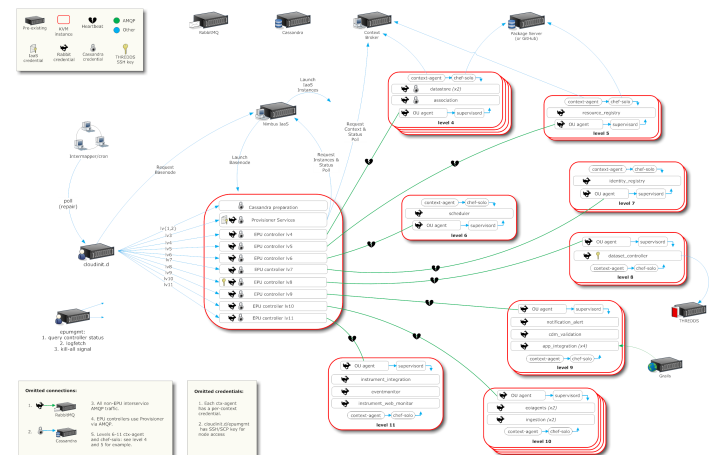
1. Each ctx-agent has a per-context credential.
2. cloudinit.d/epumgmt has SSH/SCP key for node access

- epumgmt:
1. query controller status
 2. logfetch
 3. kill-all signal

A Grid in Your Pocket...



Tim



EC2

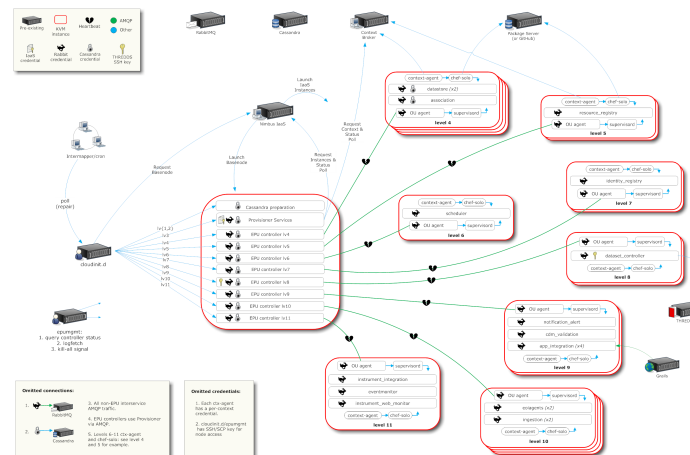
A Grid in Your Pocket...



Jamie



Tim



OOI private cloud

EC2

A Grid in Your Pocket...



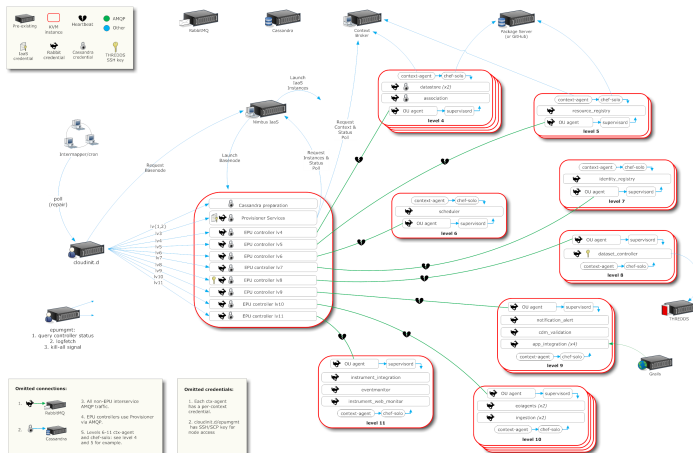
David



Jamie



Tim



FutureGrid

OOI private cloud

EC2

Cloudinit.d


Launch plan



- Repeatability: write a launch plan once, deploy many times

Cloudinit.d



Launch plan



- Deploy on cloud and non-cloud resources from many providers

Cloudinit.d


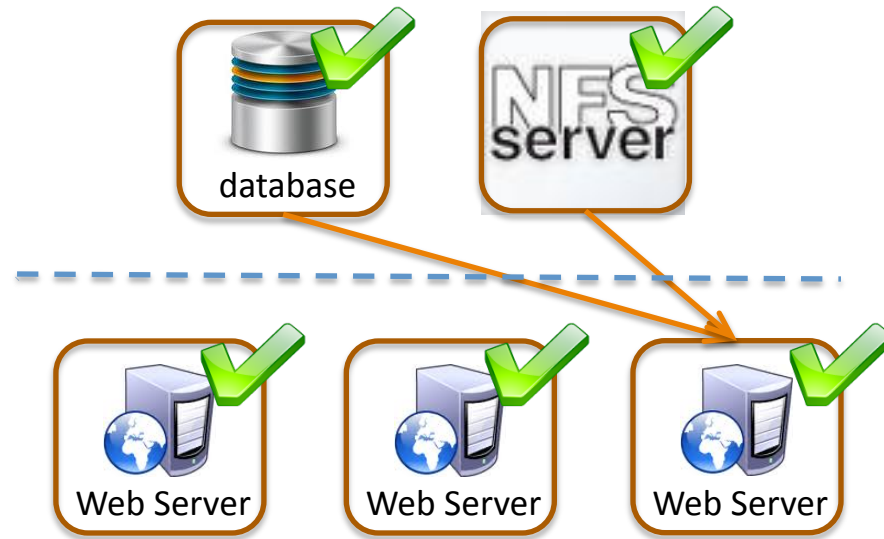
Launch plan



- Coordination of interdependent launches

Cloudinit.d

Launch plan



- User-defined launch tests

Cloudinit.d

Launch plan



database



NFS
server



Web Server




Web Server



Web Server

Run-level 1



Run-level 2



- Test-based monitoring and repair

Cloudinit.d

Launch plan



- Test-based monitoring and repair

Clouddinit.d Goals: Easy to Use

- Two roles:
 - Launch plan developers and users
- Launch plan end-users/operators
 - Lightweight
 - Copy launch plan and “one click” action
- Launch plan/application developers:
 - Minimal software assumptions (ssh)
 - “Stem cell” deployment approach
 - Incremental launch plan development

Building a Power Adapter




Elastic Scaling: Early Work

Elasticity and reliability are different sides of the same coin.

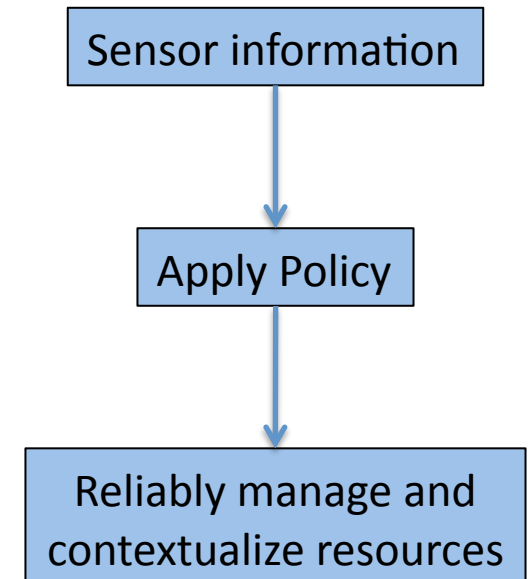
- 2008: The ALICE proof-of-concept
- 2009: ElasticSite prototype
- 2009: OOI pilot

*Challenge:
a generic HA
elastic service model*

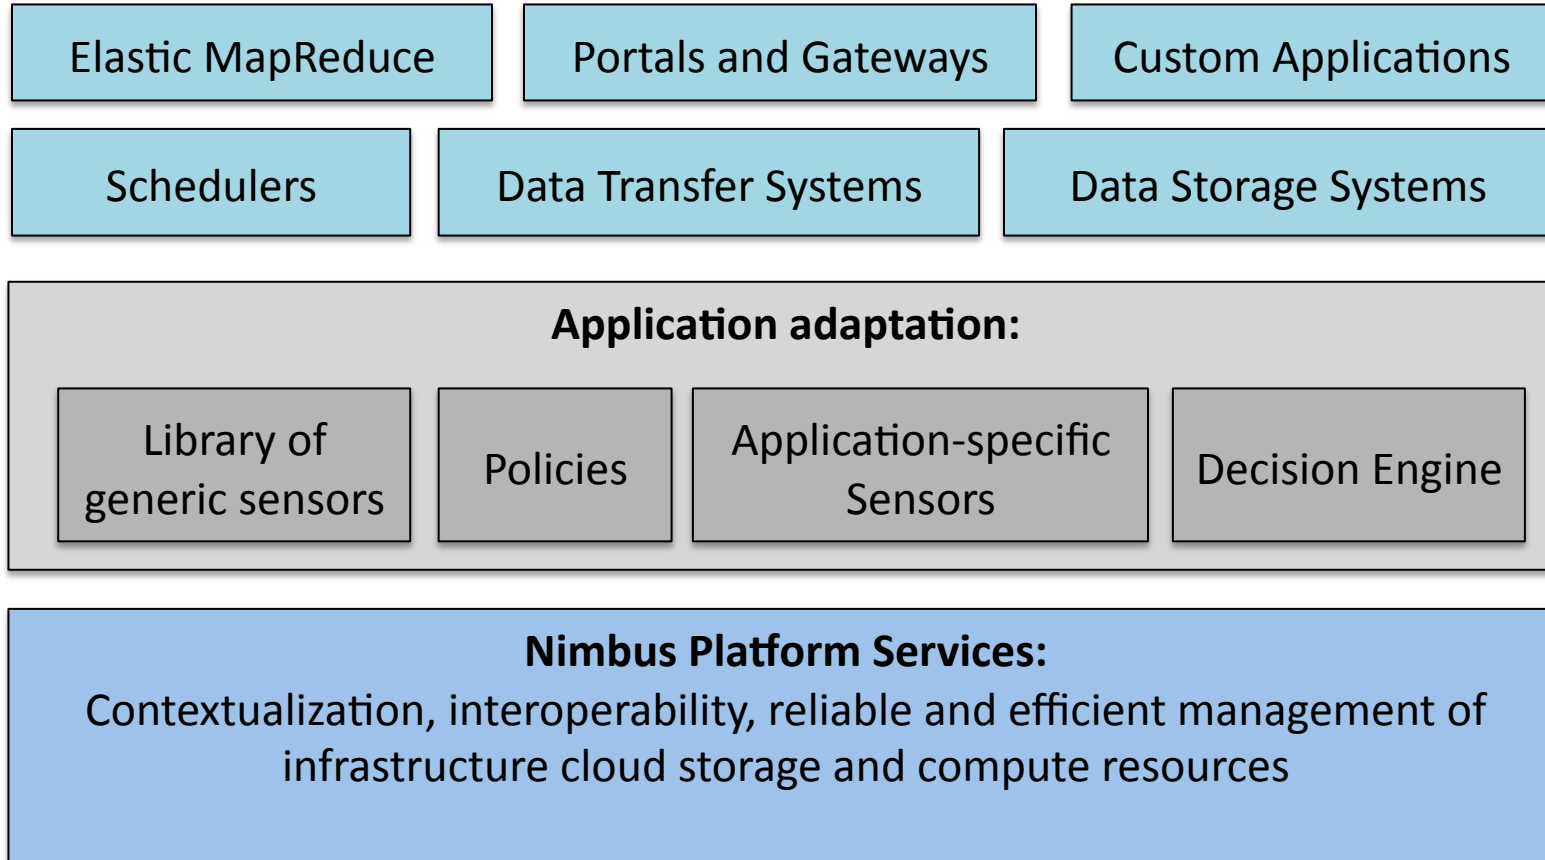


Scaling Patterns

- Monitor scaling factors:
 - A workload queue: ALiEn, PBS, AMQP,...
 - Application-specific qualities
 - Generic: deployment status, load, bank account, etc.
- Evaluate against policies
- Scale to demand
 - Across different cloud providers
 - Use contextualization to integrate machines across hybrid clouds
 - Address scientific needs for e.g., data transfer and storage
- ... then add High Availability (HA)
- Beta release later this year
 - Customizable to input, policy, decision engine, provider, etc.



Applications



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/>

Parting Thoughts

- Cloud computing is a solution!
- Computation outsourcing for science
 - Ease of use and automation
 - Science has unique outsourcing needs
 - Data management
 - A naturally bursty character
 - Escalation Pattern
- We can't have unlimited cycles...
- ...but how they are provisioned should be driven by need rather than technology limitations
- Infrastructure clouds provide such mechanism – but they are still developing
 - Security, performance, reliability, price and markets, etc.
- “We have to think anyway, we might as well think big”



www.nimbusproject.com

**Let's make cloud computing for
science happen.**