SC12 Cloud Computing for Science Tutorial: Introduction to Infrastructure Clouds

John Breshnahan, Patrick Armstrong, ***Kate Keahey****, Pierre Riteau*
Argonne National Laboratory
Computation Institute, University of Chicago

Outsourcing Computing for Science

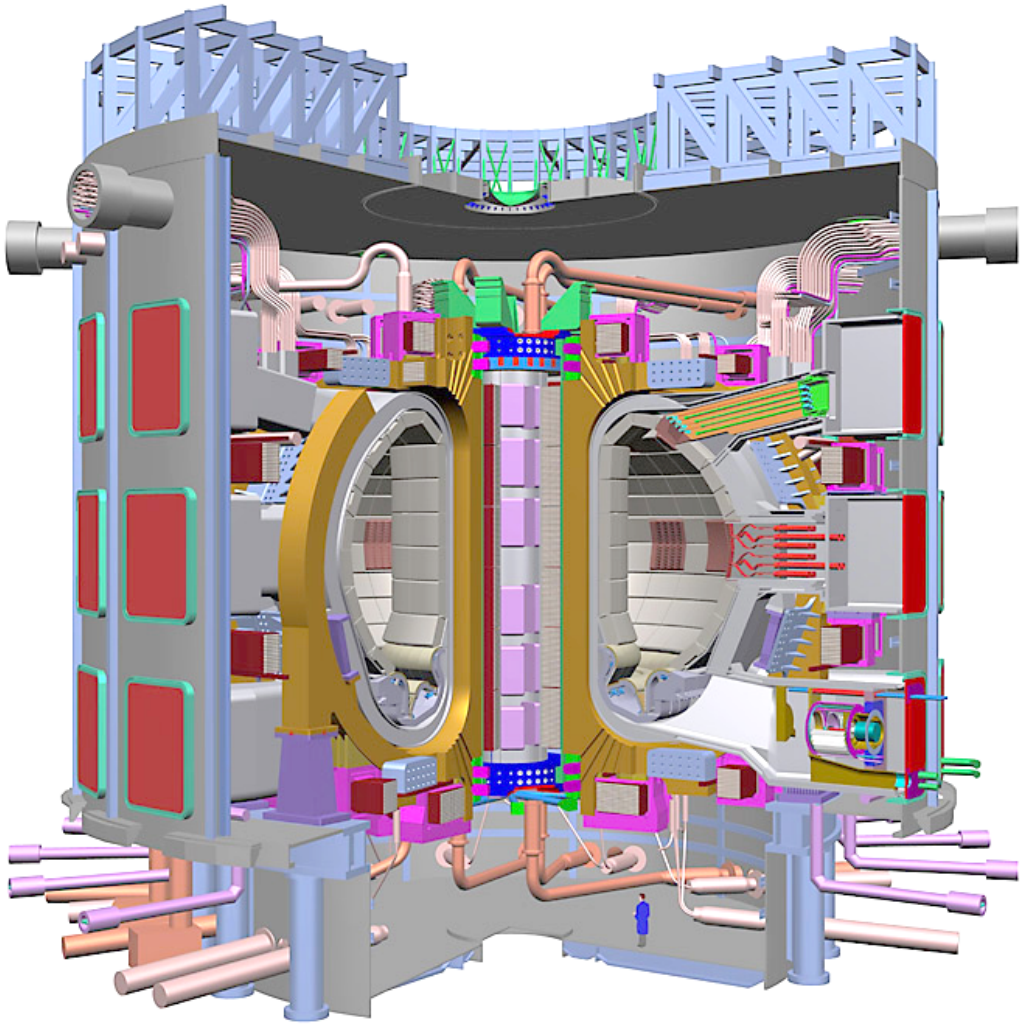

Control over environment

- Complexity
- Consistency

Control over availability

- Many real-time scenarios
- Naturally bursty character

NIMBUS www.nimbusproject.org

Introducing Infrastructure Clouds

- Infrastructure-as-a-Service (IaaS) Cloud Computing
  - Access to remote storage and compute resources
- Benefits:
  - On-demand access
  - Pay-as-you-go charging model
  - Virtualization, environment control
  - Convenience via outsourcing

Everything as a Service

Software-as-a-Service (SaaS)

Community-specific tools, applications and portals

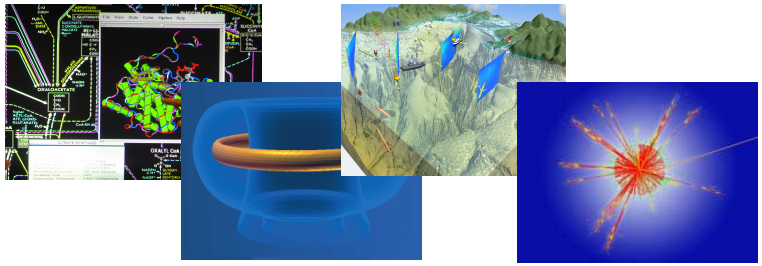

Platform-as-a-Service (PaaS)

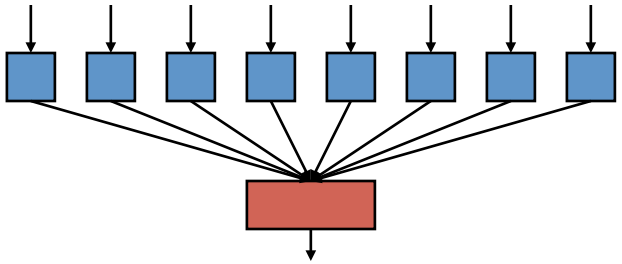

Microsoft

Google

Infrastructure-as-a-Service (IaaS)

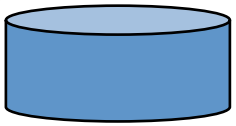

amazon web services™

flexiscale™

GOGRID beta
A ServePath Company

NIMBUS

www.nimbusproject.org

Control

Specialization

Infrastructure-as-a-Service: Quickstart

IaaS Modus Operandi

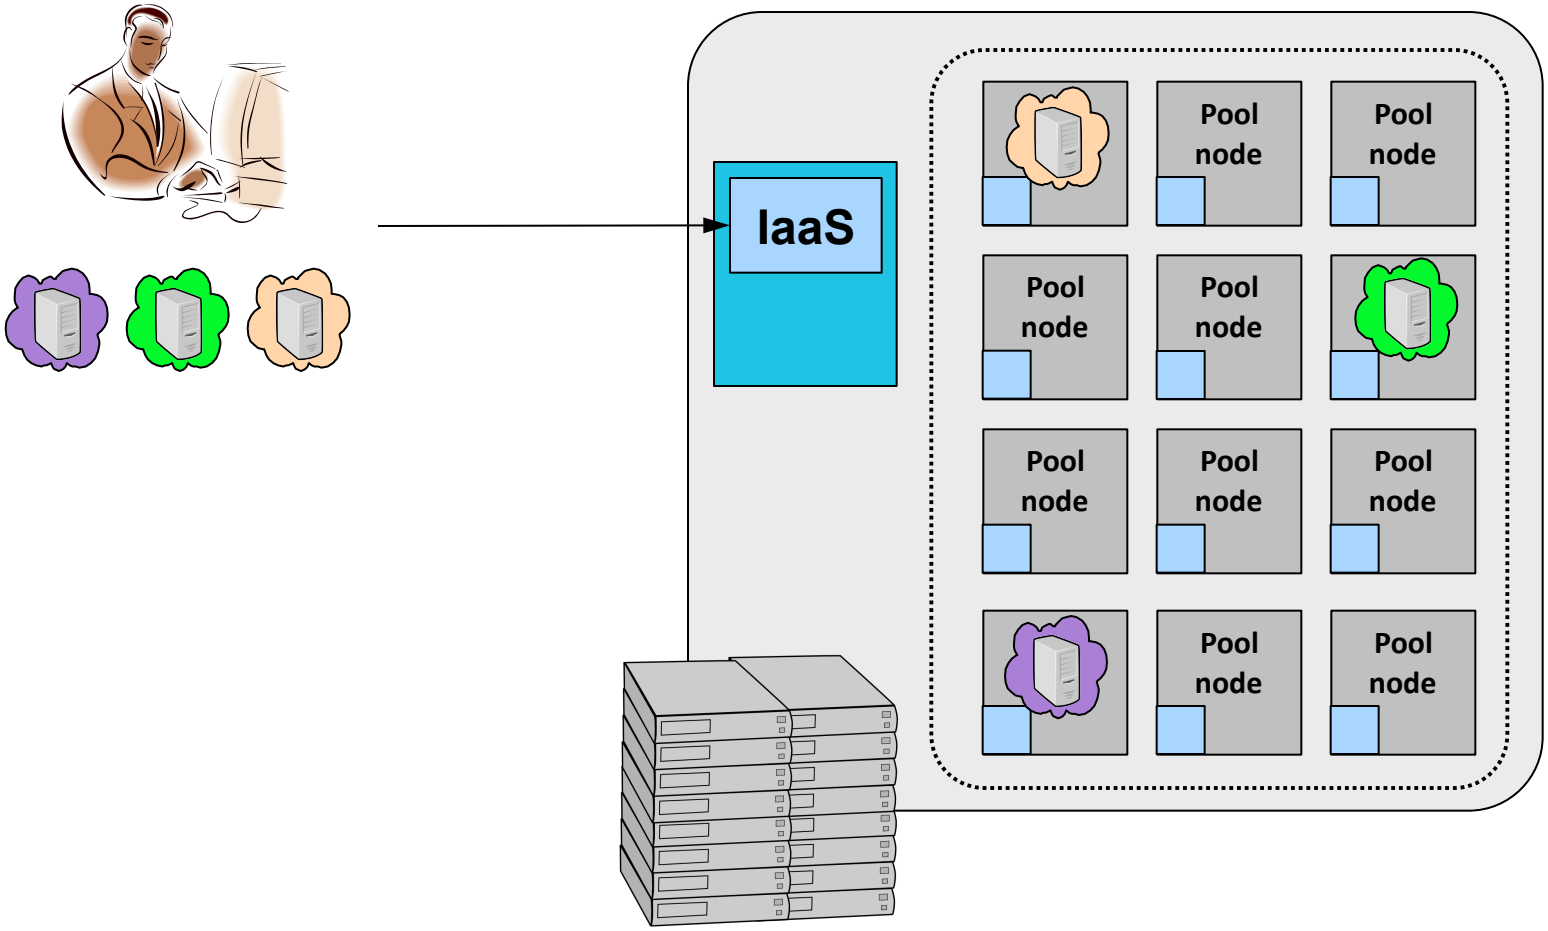

IaaS Modus Operandi

The IaaS service publishes information about each VM

Users can find out information about their VM (e.g. what IP the VM was bound to)

Users can interact directly with their VM in the same way the would with a physical machine (e.g., ssh).

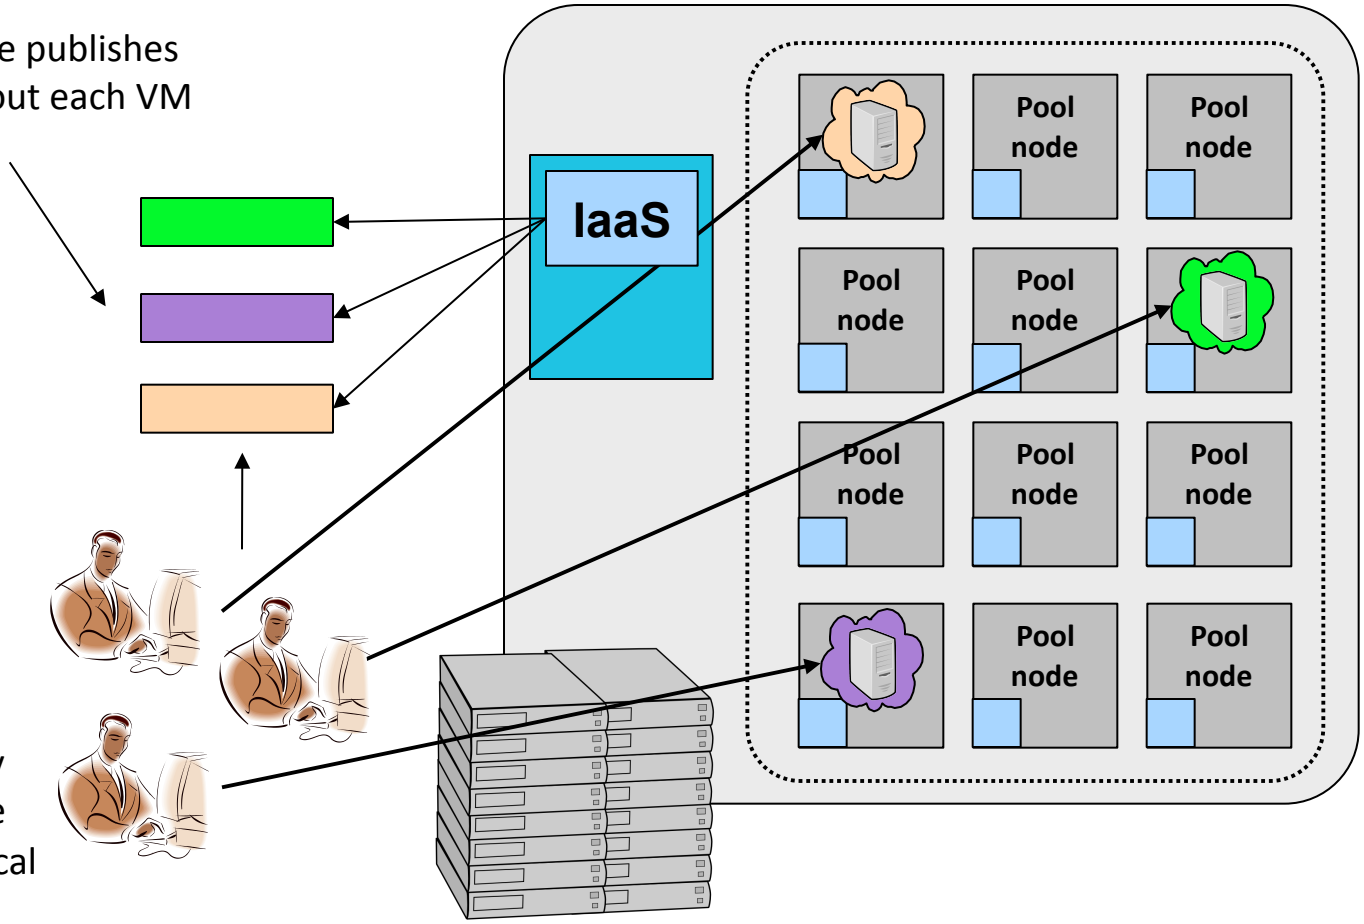

Types of IaaS Resources

- Resource shapes/types
  - Bundles of virtual resource parameters
  - Exact (memory/storage) and vague (I/O performance, “compute units”)
  - Special hardware (e.g., GPUs)
  - Different types of storage options: e.g., S3 vs EBS
- Resource availability/persistence
  - On-demand instances
  - Subscription instances (“reserved” instance)
  - **Spot instances**
  - Standard vs reduced redundancy
- Pricing models
  - From 2 cents to ~$3 per hour for on-demand instances
  - Consolidated billing
  - Storage: per storage, access, and outgoing transfer

Infrastructure Cloud Resources

Community clouds

sciencecloud.org

Magellan
a cloud for science

… also various MRI projects, WestGrid, Grid'5000

Commercial clouds

amazon web services™

rackspace HOSTING

Microsoft

flexiscale™

GOGRID beta
A ServePath Company

Google

Configure your own private cloud

NIMBUS

openstack™

Eucalyptus

IaaS Modus Operandi

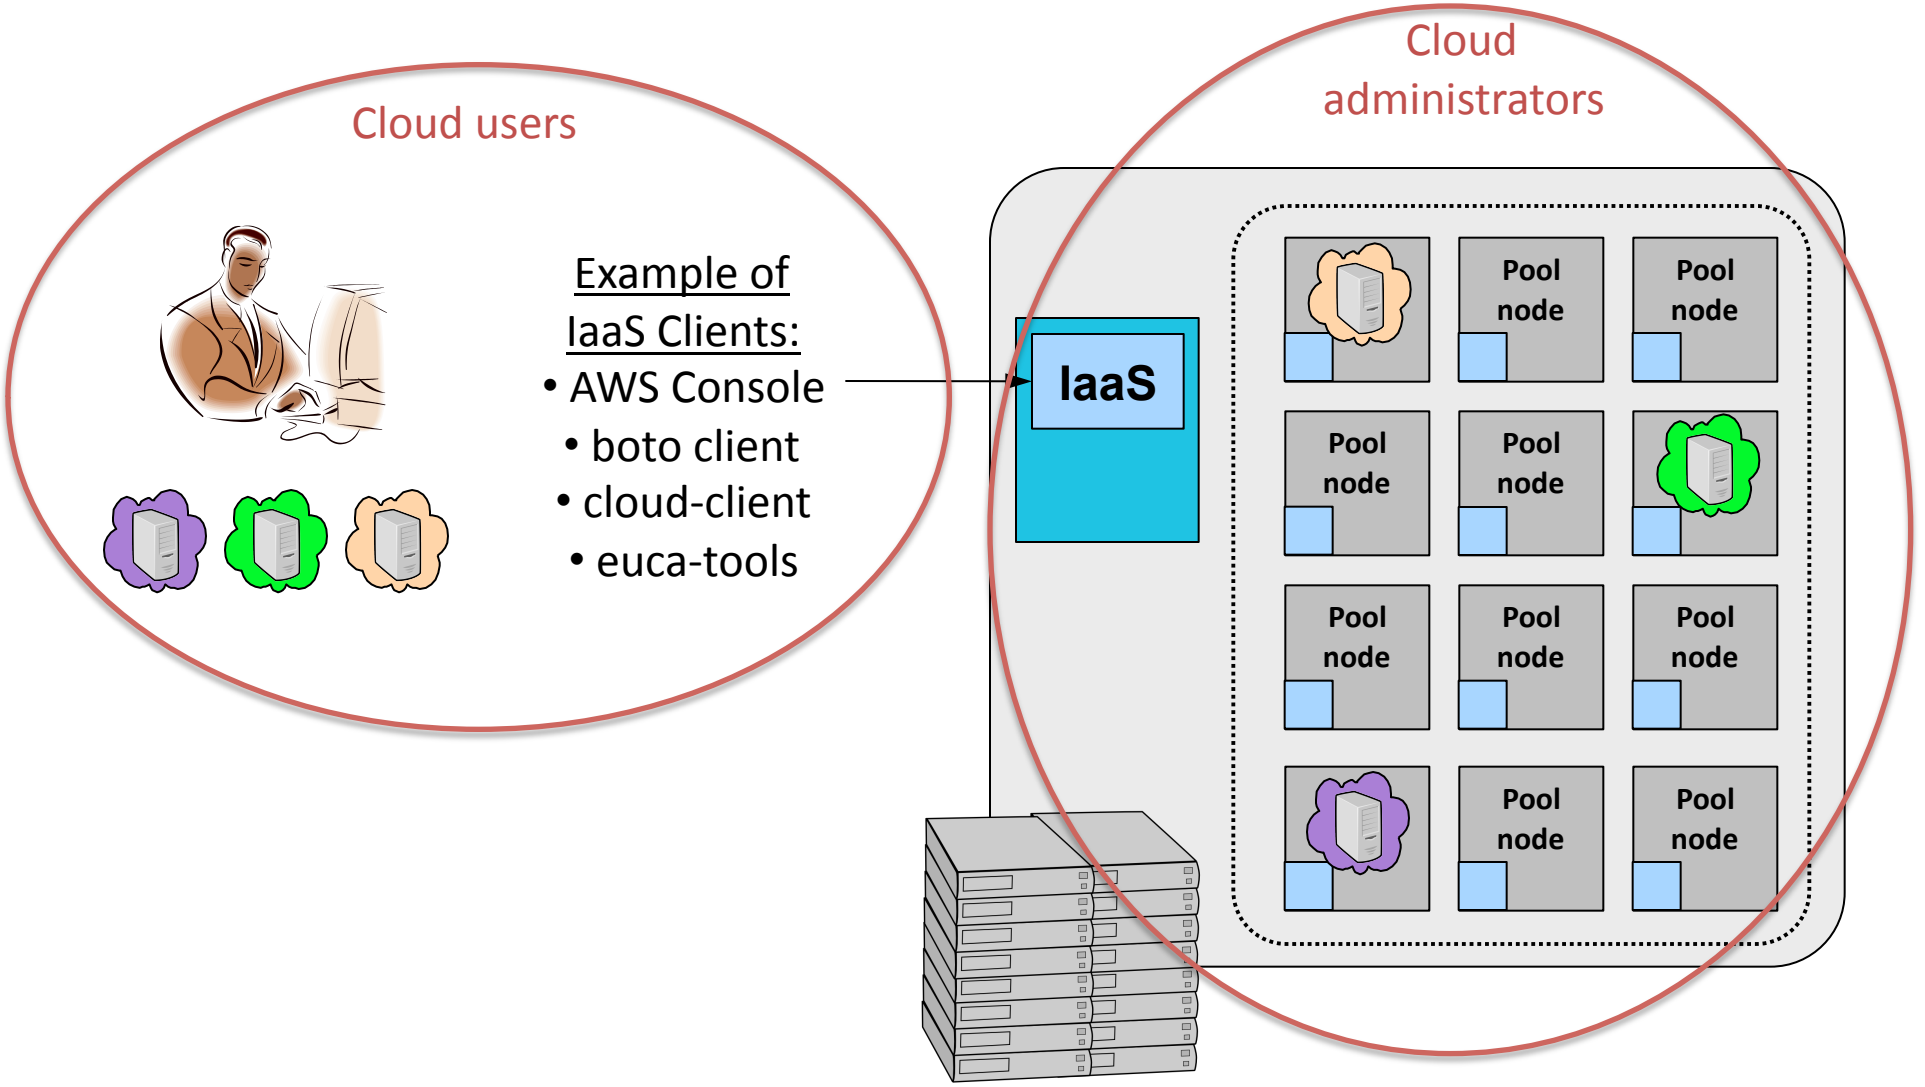

Applications:
What Can You Do with Infrastructure Clouds?

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
- Approach:
  - Running on user's desktop
  - Then running on Nimbus Science Clouds, Magellan and EC2
  - A platform for building appliances representing push-button pipelines
  - Leveraging spot instances
- Impact
  - From desktop to cloud
  - http://clovr.org

Edition Comparison

	Skeleton	Base	Standard
Ubuntu 10.04	✔	✔	✔
Grid Engine	✖	✔	✔
Hadoop	✖	✔	✔
Ganglia	✖	✔	✔
Vappio	✖	✔	✔
Ergatis	✖	✖	✔
Platforms			
EC2			
Eucalyptus			
VirtualBox			
VMware			
Xen			
Magellan Cloud			
Science Clouds			

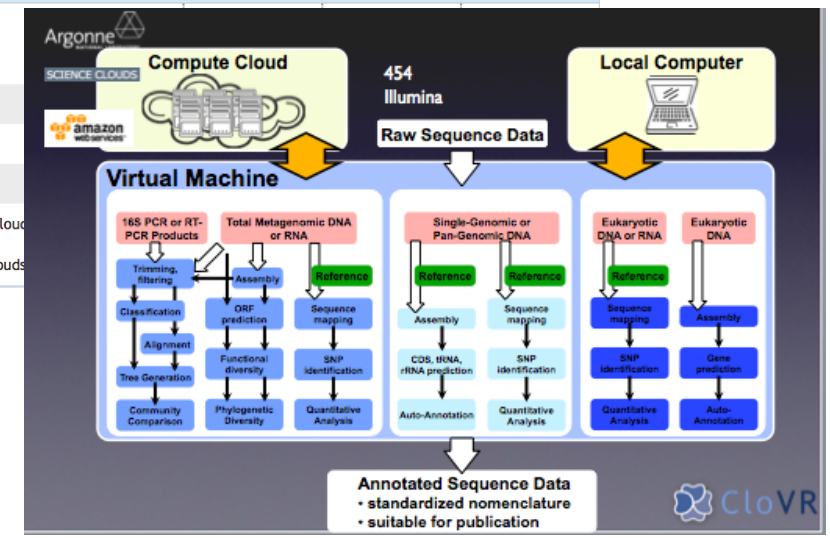

CANFAR

Canadian Advanced Network for Astronomical Research

University of Victoria NRC·CNRC

Work by the UVIC team

- 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

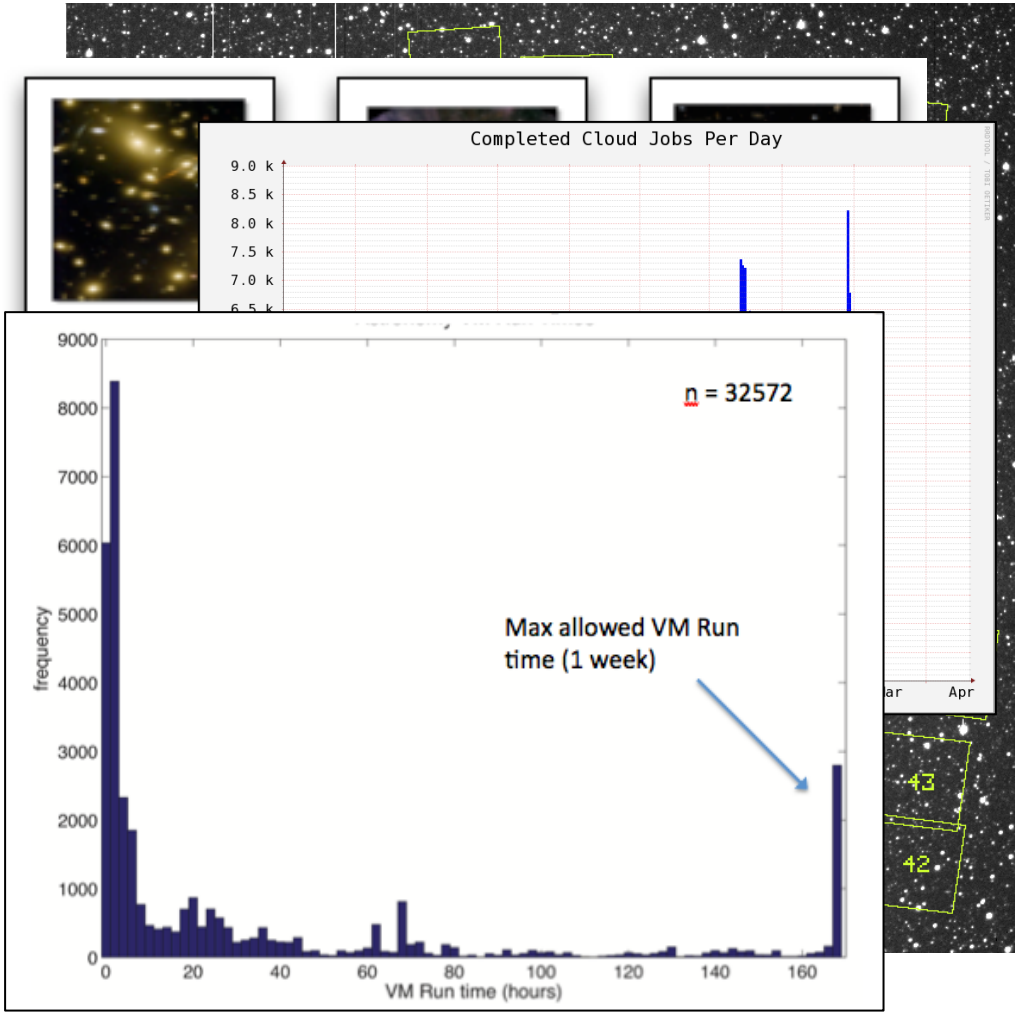

Sky Computing

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

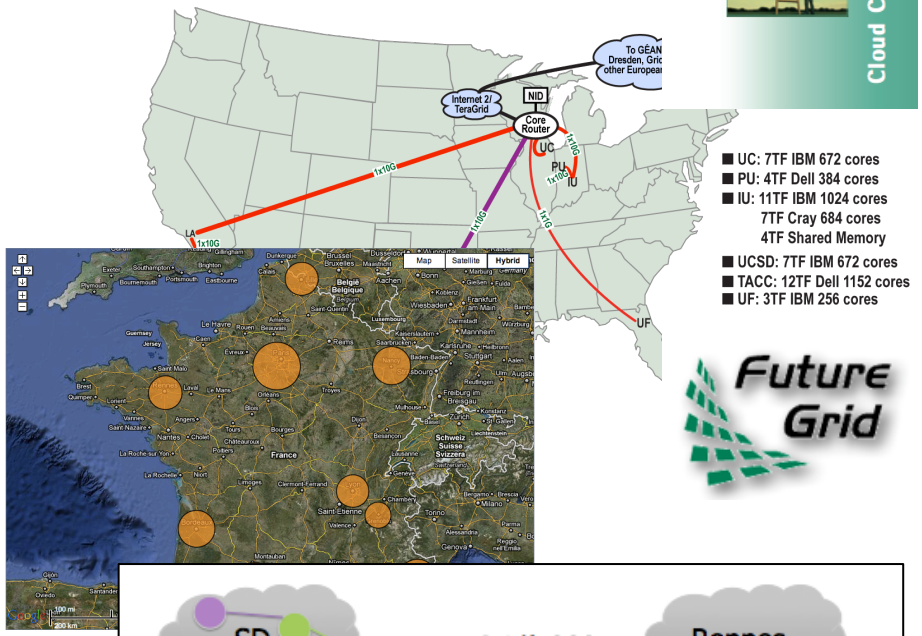

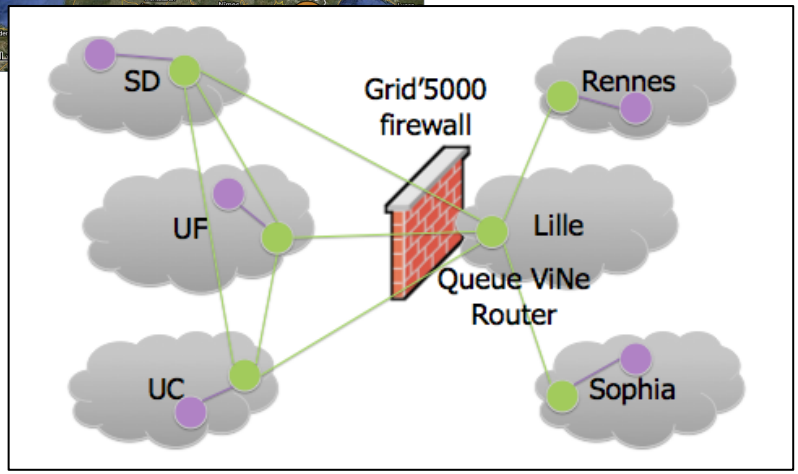

BaBar

Canadian Efforts

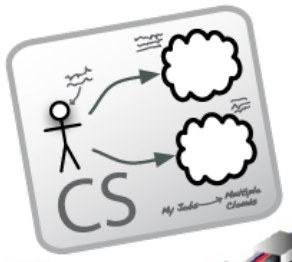

University of Victoria NRC·CNRC

Work by the UVIC team

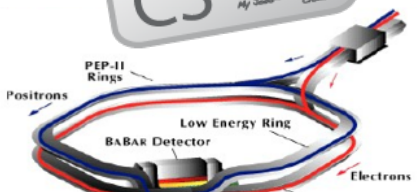

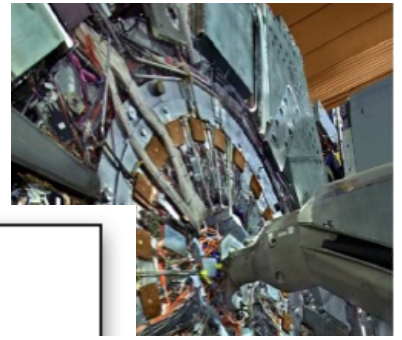

- BarBar Experiment at SLAC in Stanford, CA
- Using clouds to simulate electron-positron collisions in their detector
- Exploring virtualization as a vehicle for data preservation
- Approach:
  - Appliance preparation and management
  - Distributed Nimbus clouds
  - Cloud Scheduler
- Running production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

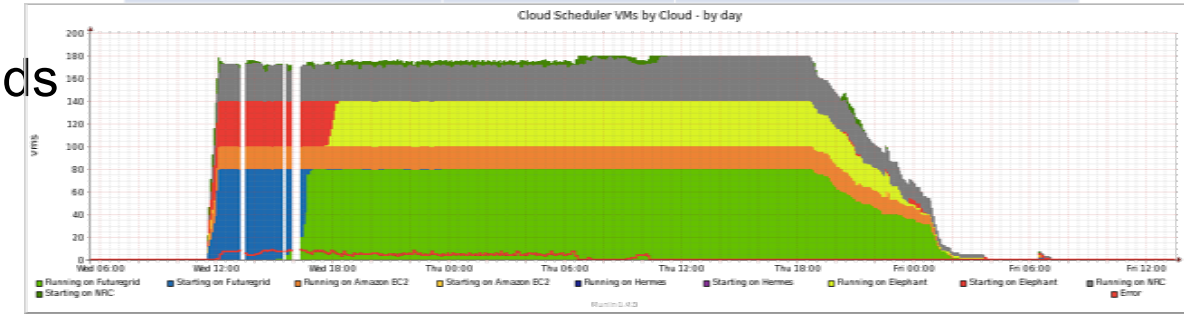

Ocean Observatory Initiative

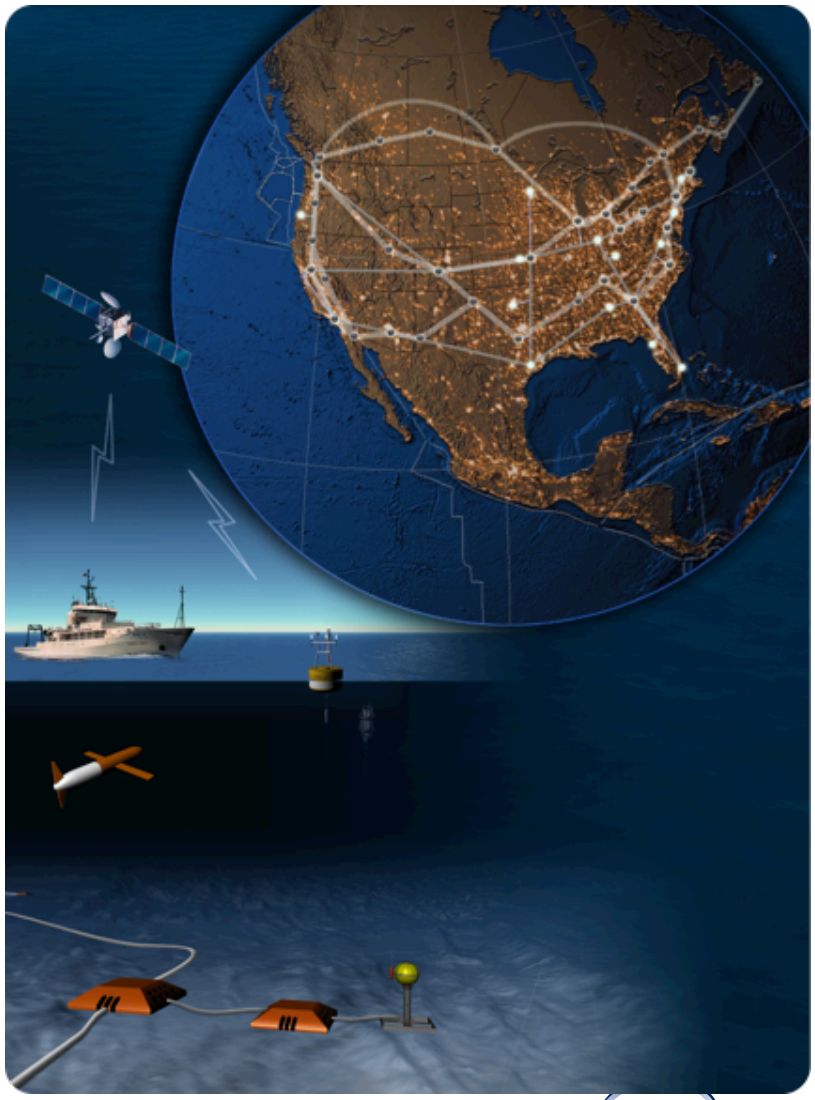

- Towards Observatory Science
- Sensor-driven processing
  - Real-time event-based data stream processing capabilities
  - Highly volatile need for data distribution and processing
  - An “always-on” service
- Nimbus team building platform services for integrated, repeatable support for on-demand science
  - High-availability
  - Auto-scaling
- From regional Nimbus clouds to commercial clouds

The STAR Experiment

- Use case: cloud processing for W-boson reconstruction
- Overall achievement
  - "10 months became 3 months"
- Benefits:
  - Reduce "time to science"
  - Near real-time processing

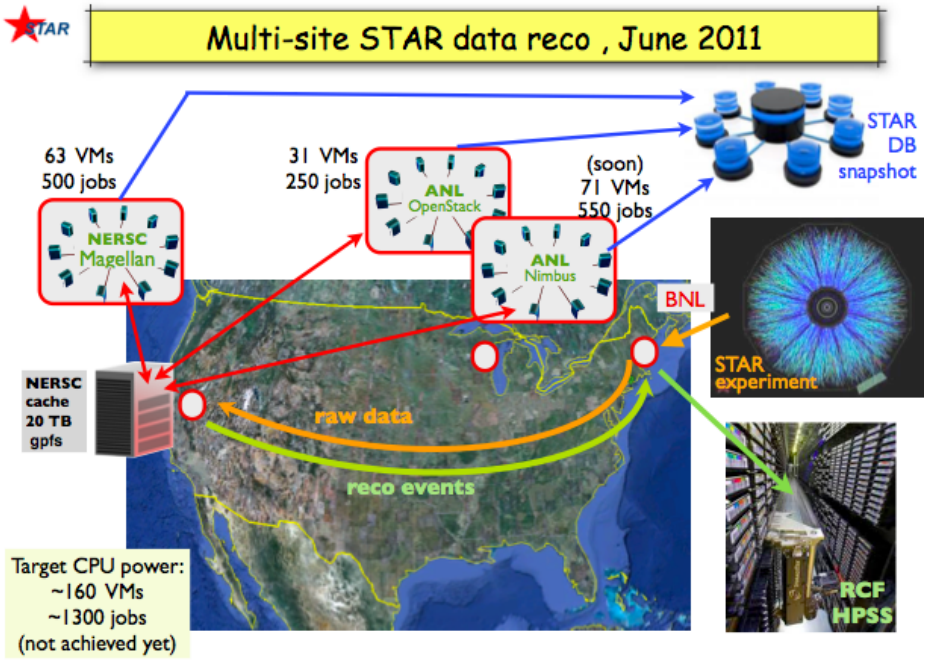

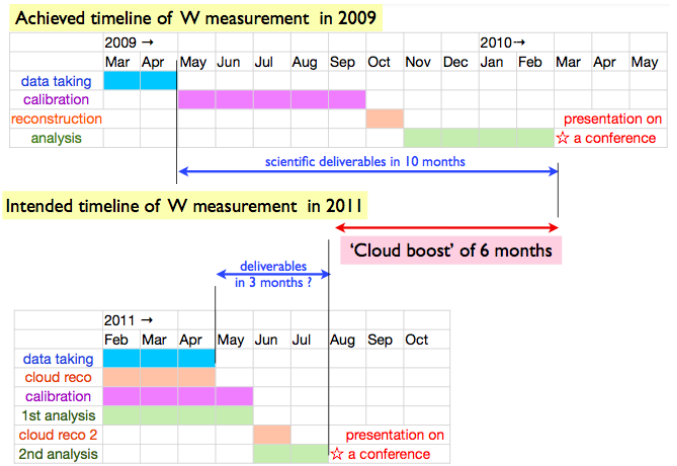

Trends and Patterns

- On-demand, **elastic** processing…
  - Observatories, experiments, conference deadlines, fluctuating workload, growth management…
- … over multiple providers….
  - **Risk-mitigation:** not enough cycles, failure, market factors
- … with elementary ease of use…
  - **Automated** provisioning of infrastructure resources
  - Harness the power while avoiding complexity
- …facilitates going from one-offs to **production** runs…
  - Steadily increasing in both size and buy-in
- …given the right model
  - **Failure-prone** environment
  - Management at scale
  - Achieving "uninterrupted power supply"

Infrastructure Platform Requirements

Turnkey Virtual Clusters

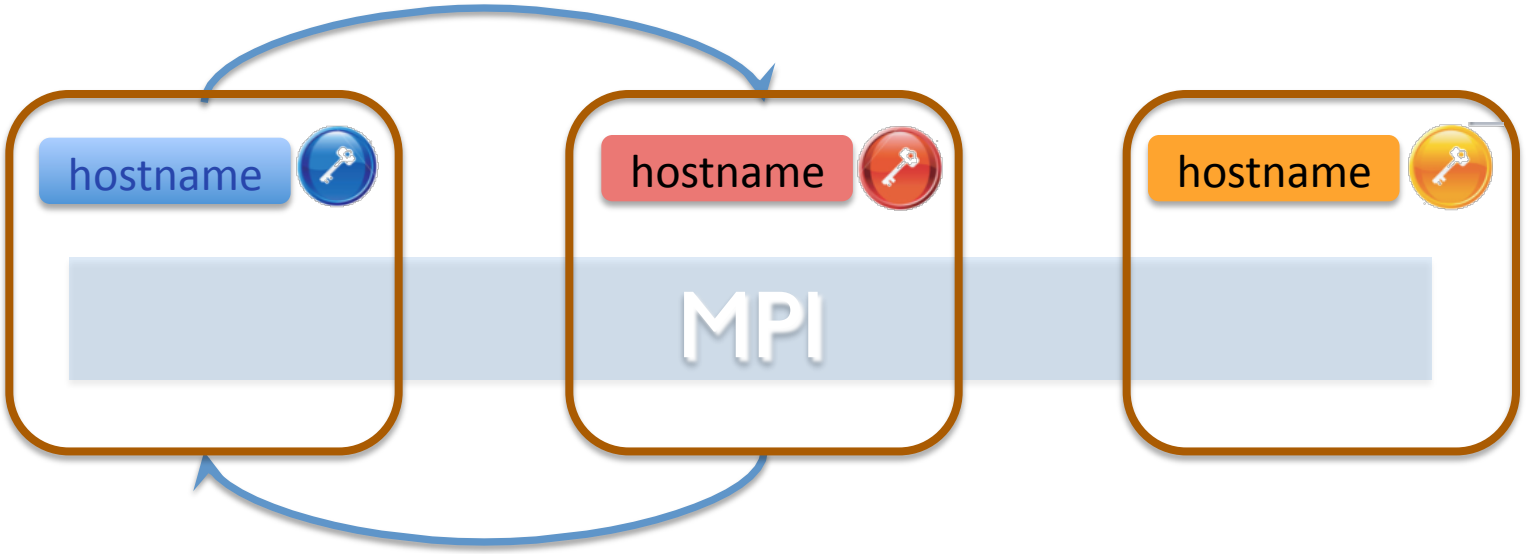

- Configuring a turnkey virtual cluster
  - Shared trust/security context
  - Shared configuration/context information

Managing Complexity

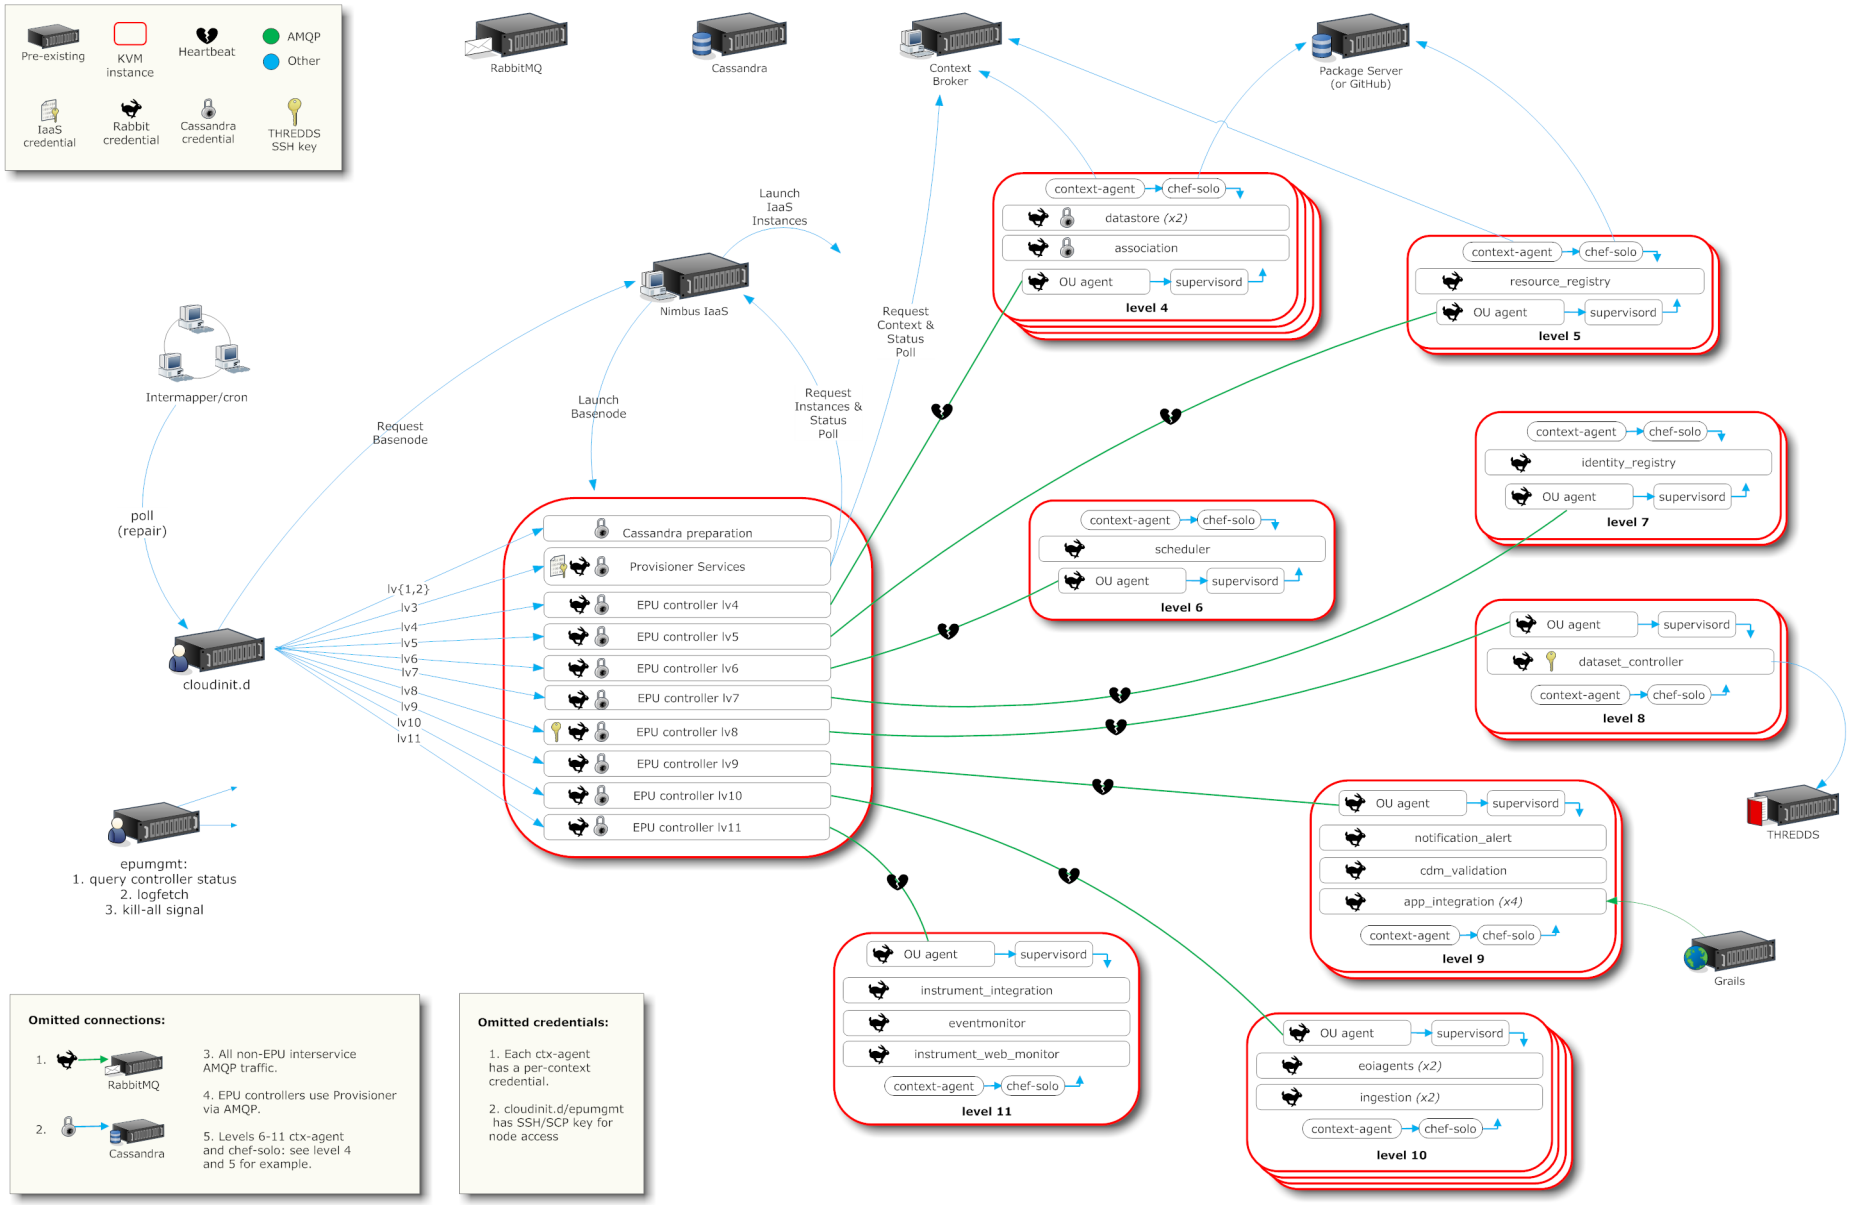

Elastic Scaling Examples

- The ALICE Example
- ElasticSite prototype
- OOI pilot

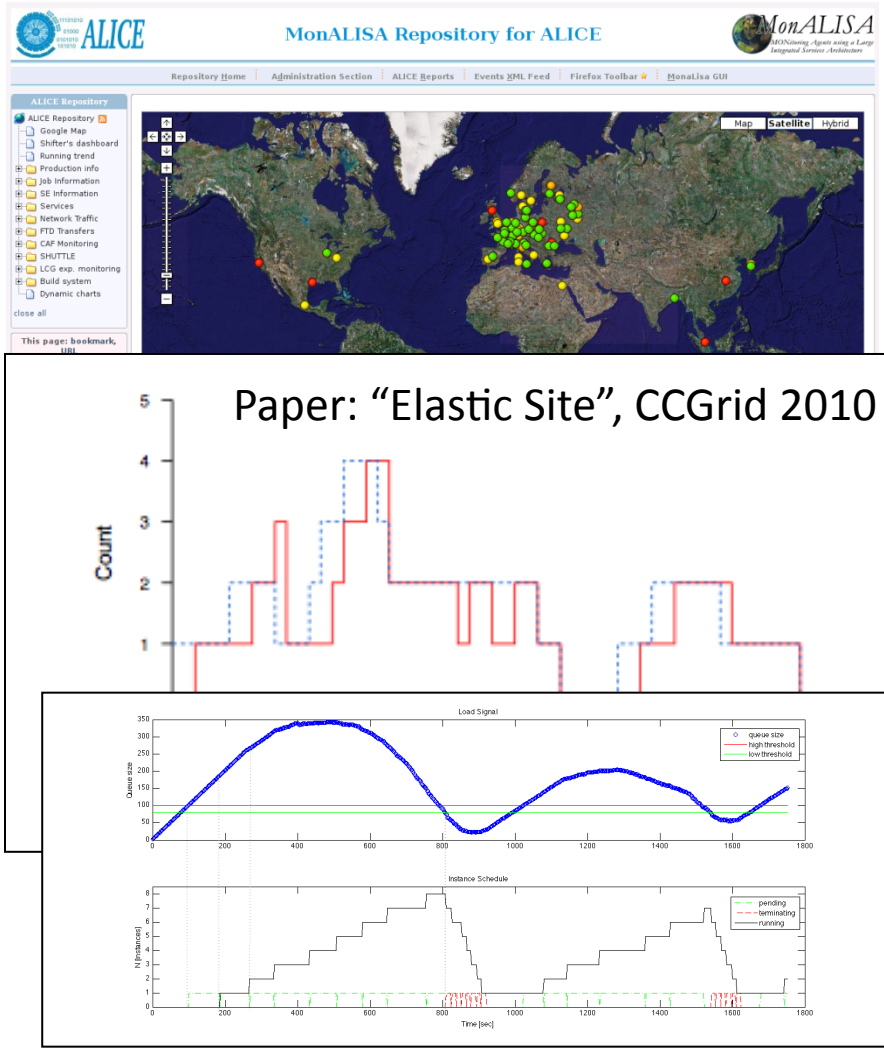

Plan for the Day

- First things first: start using IaaS clouds!
  - AWS demonstration
  - Do it on your own: deploy a VM on FutureGrid!
- Which cloud should I use?
  - Comparison of commercial and community clouds
- How can I go beyond deploying a few VMs?
  - Infrastructure Platform Tools: virtual clusters, managing deployment complexity, autoscaling and high availability
- Put it all in practice
  - Applications: environment management, bioinformatics, high-energy physics