Infrastructure-as-a-Service Cloud Computing for Science

April 2010

Salishan Conference

Kate Keahey

keahey@mcs.anl.gov

Nimbus project lead

Argonne National Laboratory

Computation Institute, University of Chicago

Argonne NATIONAL LABORATORY

NSF

Cloud Computing for Science

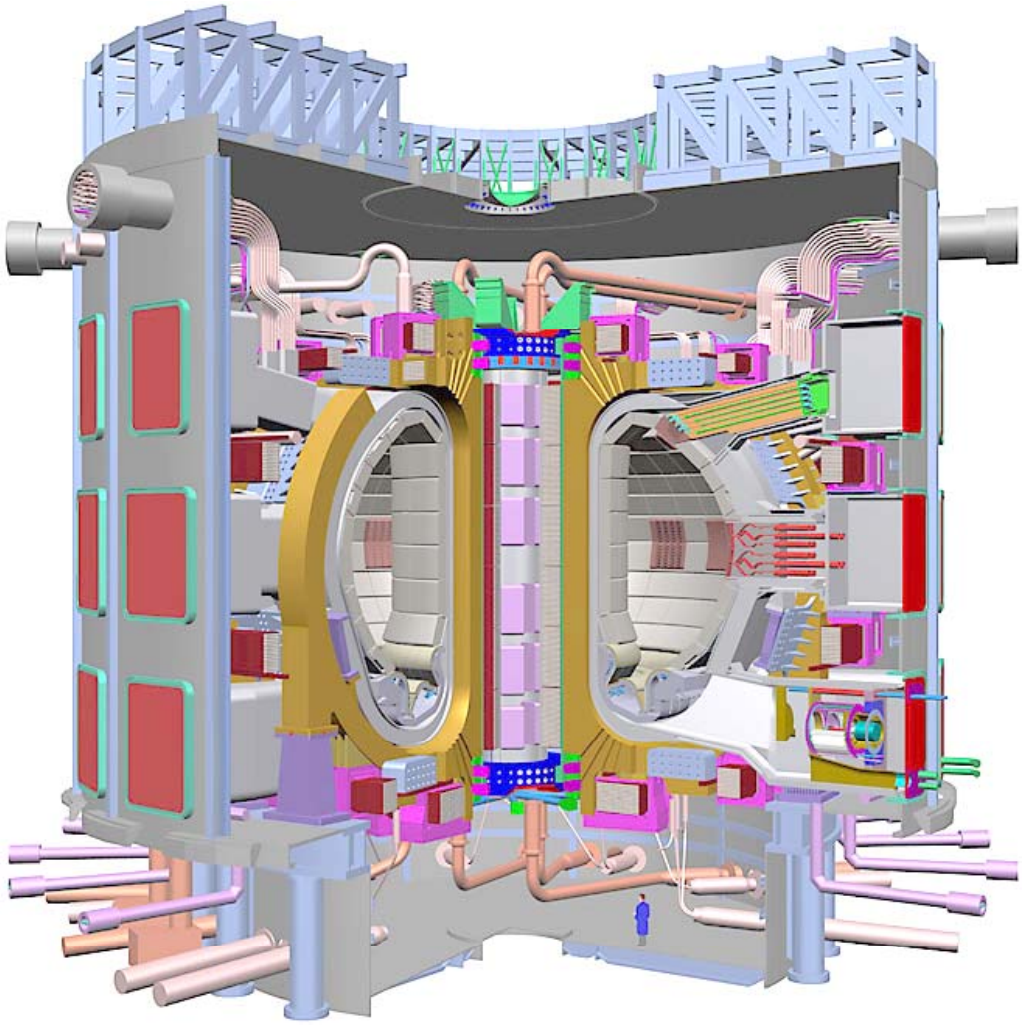

- Environment control
- Resource control

“Workspaces”

- Dynamically provisioned environments
  - Environment control
  - Resource control
- Implementations
  - Via leasing platforms: reimaging, configuration, dynamic account
  - Via virtualization

Isolation

The Nimbus Workspace Service

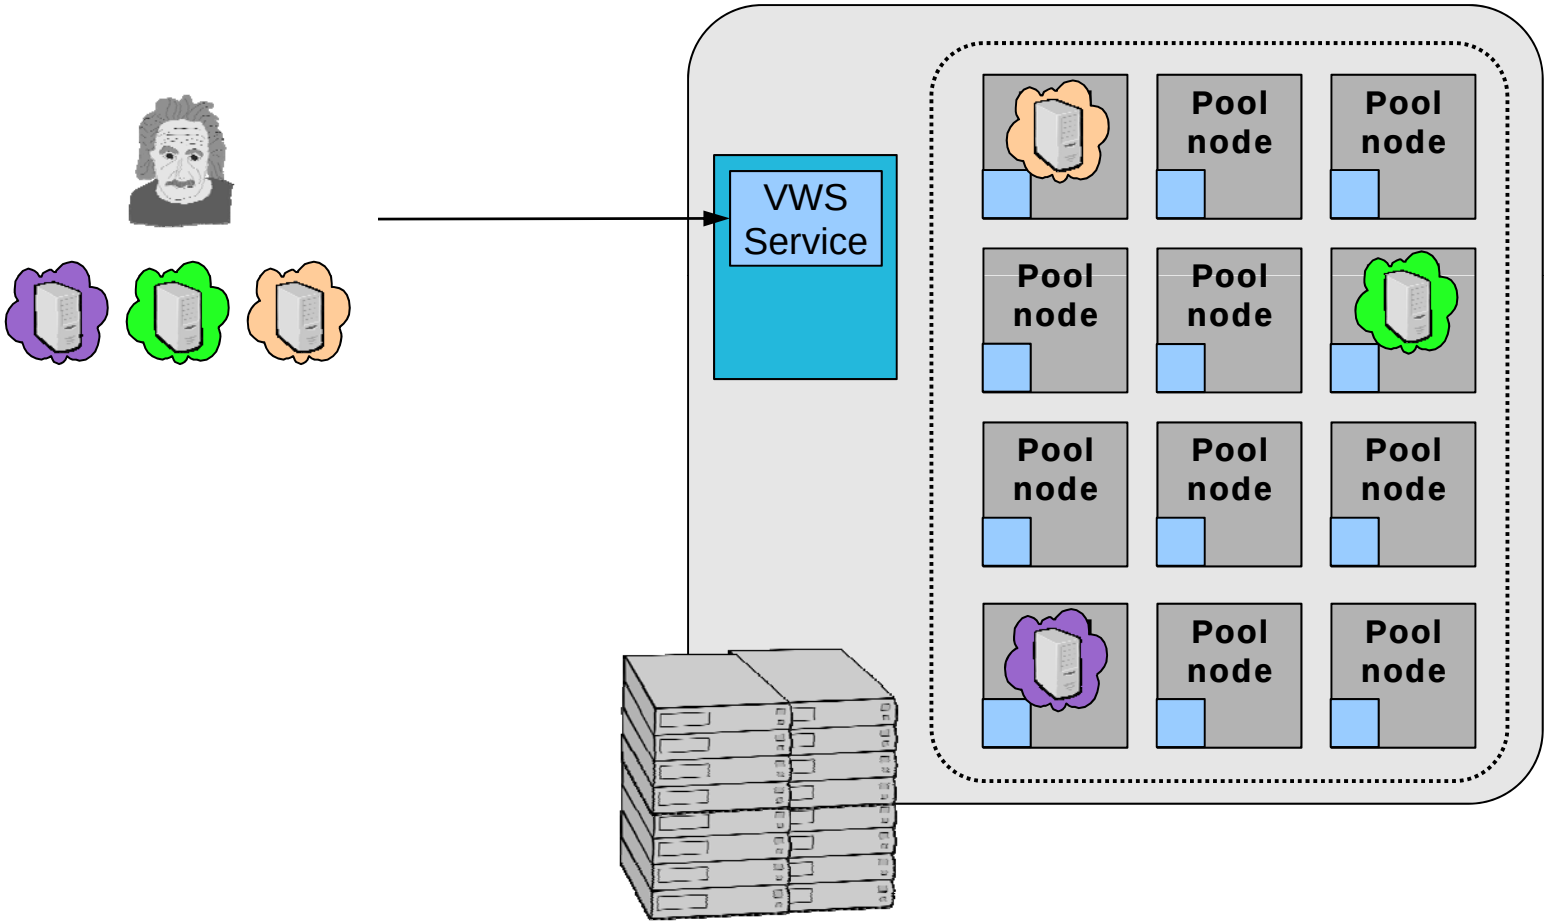

The Nimbus Workspace Service

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.

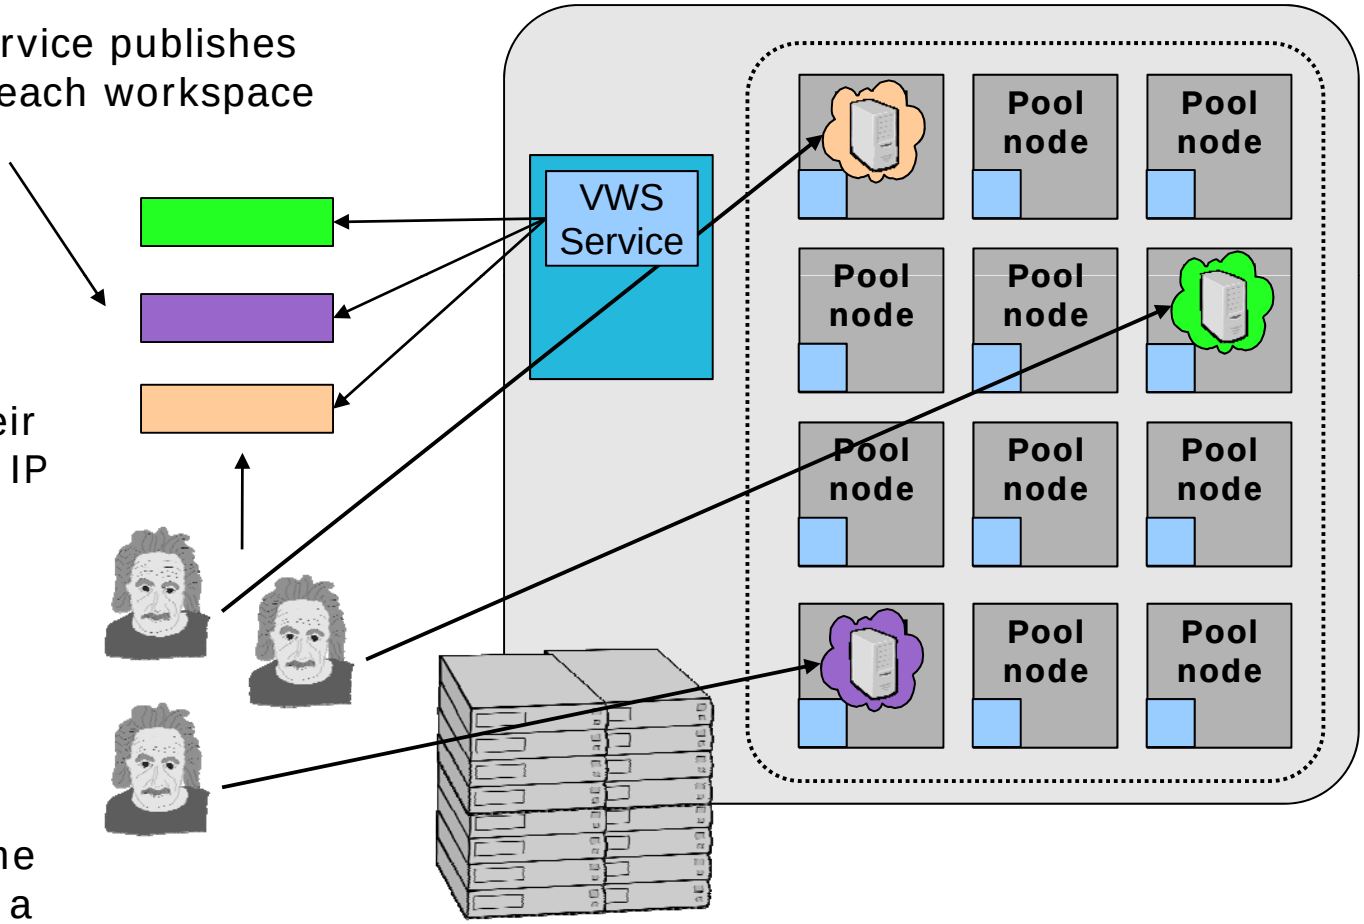

Nimbus: Cloud Computing for Science

- Allow providers to build clouds
  - Workspace Service: a service providing EC2-like functionality
  - WSRF and WS (EC2) interfaces
- Allow users to use cloud computing
  - Do whatever it takes to enable scientists to use IaaS
  - Context Broker: turnkey virtual clusters,
  - Also: protocol adapters, account managers and scaling tools
- Allow developers to experiment with Nimbus
  - For research or usability/performance improvements
  - Open source, extensible software
  - Community extensions and contributions: UVIC (monitoring), IU (EBS, research), Technical University of Vienna (privacy, research)
- Nimbus: www.nimbusproject.org

Clouds for Science: a Personal Perspective

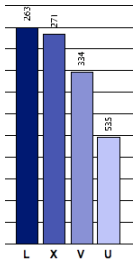

Xen released

EC2 released

STAR

First STAR production run on EC2

Science Clouds available

OOI starts

Future Grid

Magellan
a cloud for science

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

STAR experiment

Work by Jerome Lauret, Leve Hajdu, Lidia Didenko (BNL), Doug Olson (LBNL)

- STAR: a nuclear physics experiment at Brookhaven National Laboratory
- Studies fundamental properties of nuclear matter
- Problems:
  - Complexity
  - Consistency
  - Availability

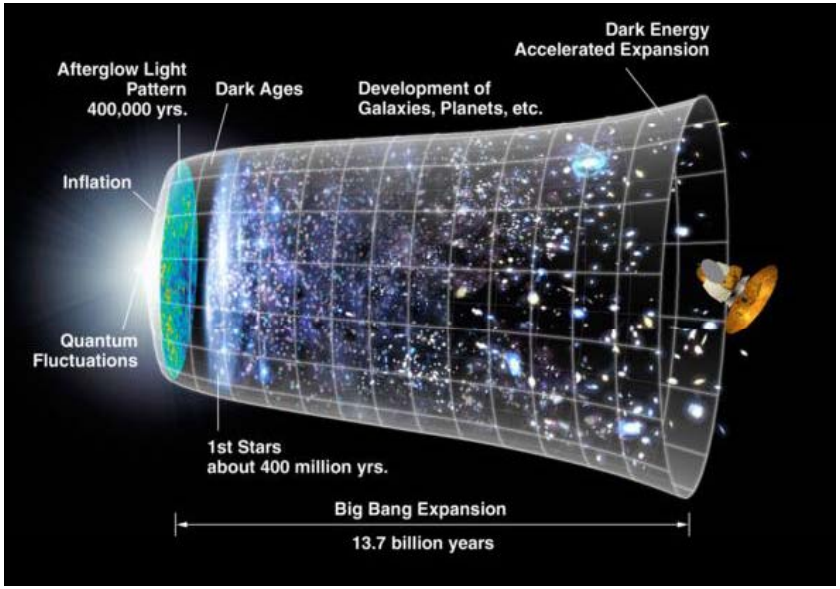

STAR Virtual Clusters

- 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
- The Quark Matter run: producing just-in-time results for a conference: http://www.isgtw.org/?pid=1001735

Priceless?

- Compute costs: $ 5,630.30
  - 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
- Storage costs: $ 4.69
  - Images only, all data transferred at run-time
- Producing the result before the deadline...

...$ 5,771.37

Cloud Bursting

ALICE: Elastically Extend a Grid

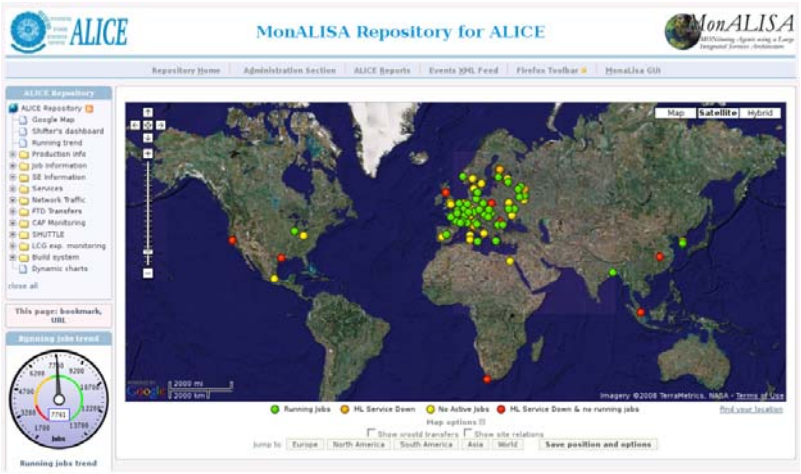

Elastically Extend a cluster

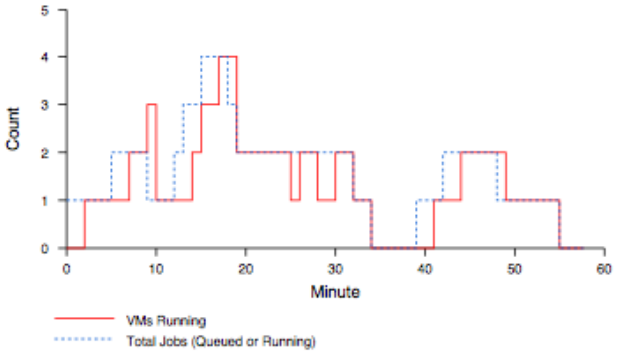

OOI: Provide a Highly Available Service

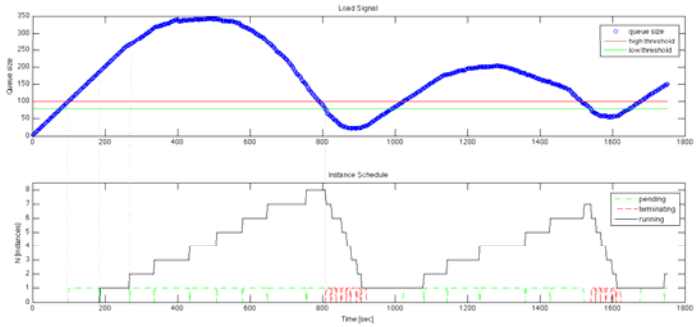

React to Emergency

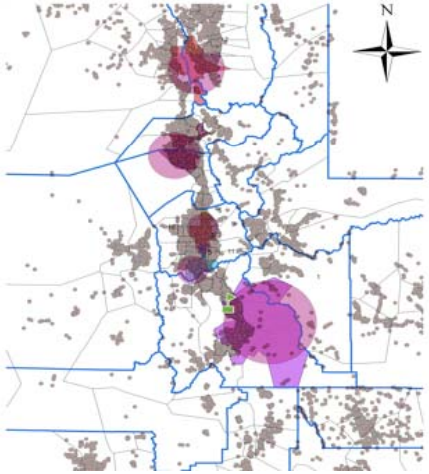

Hadoop in the Science Clouds

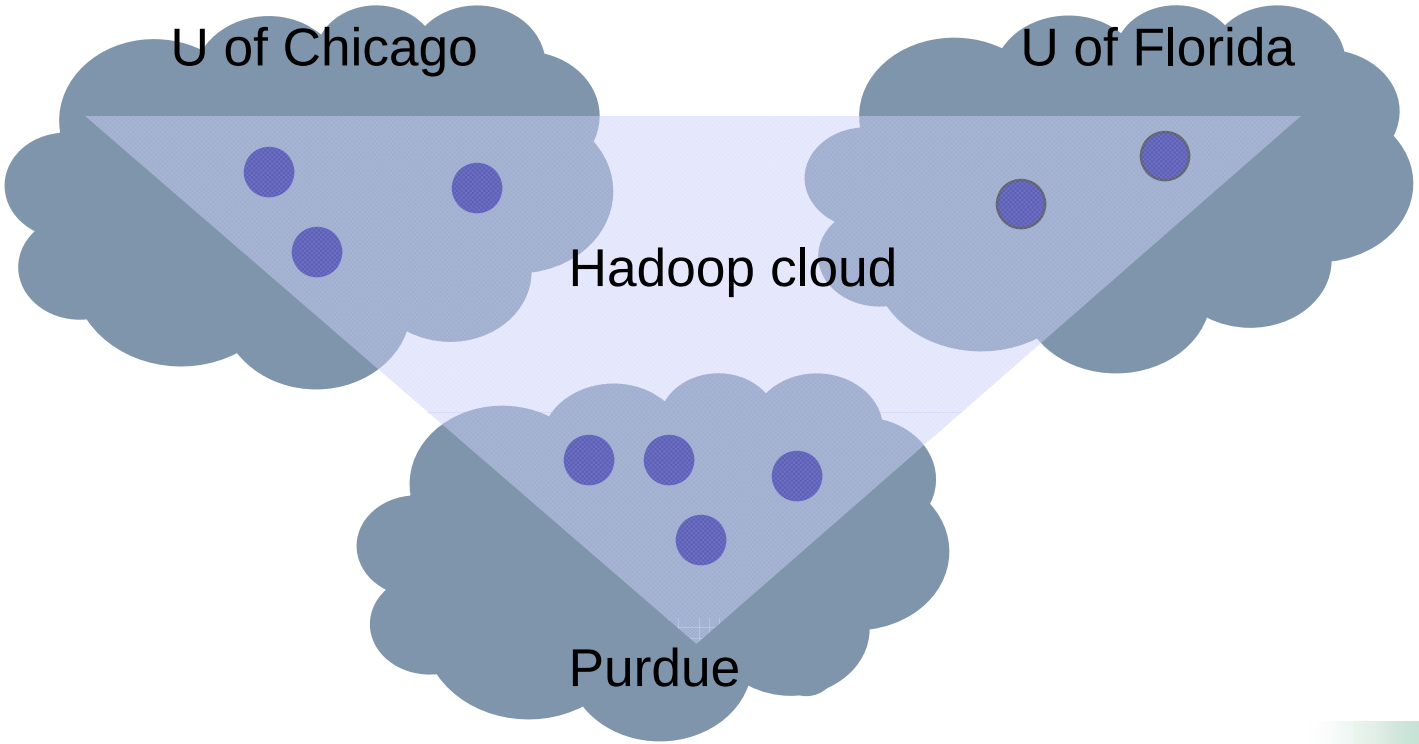

- *Papers:*
  - *"CloudBLAST: Combining MapReduce and Virtualization on Distributed Resources for Bioinformatics Applications" by A. Matsunaga, M. Tsugawa and J. Fortes. eScience 2008.*
  - *"Sky Computing", by K. Keahey, A. Matsunaga, M. Tsugawa, J. Fortes, to appear in IEEE Internet Computing, September 2009*

Genomics: Dramatic Growth in the Need for Processing

- moving from big science (at centers) to ***many*** players
- democratization of sequencing
- we currently have more data than we can handle

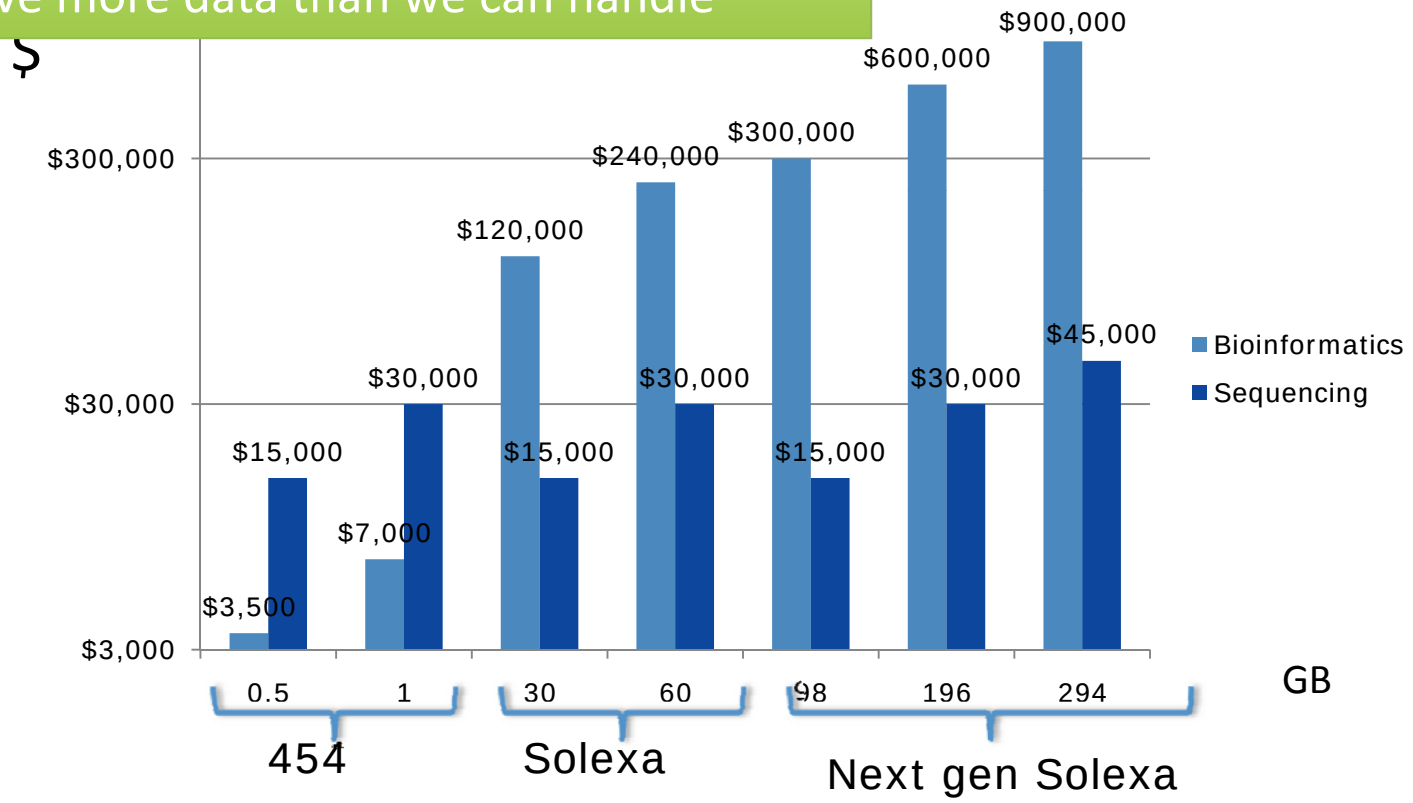

From *Folker Meyer, "The M5 Platform"*

Ocean Observatory Initiative

CI: Linking the marine infrastructure to science and users

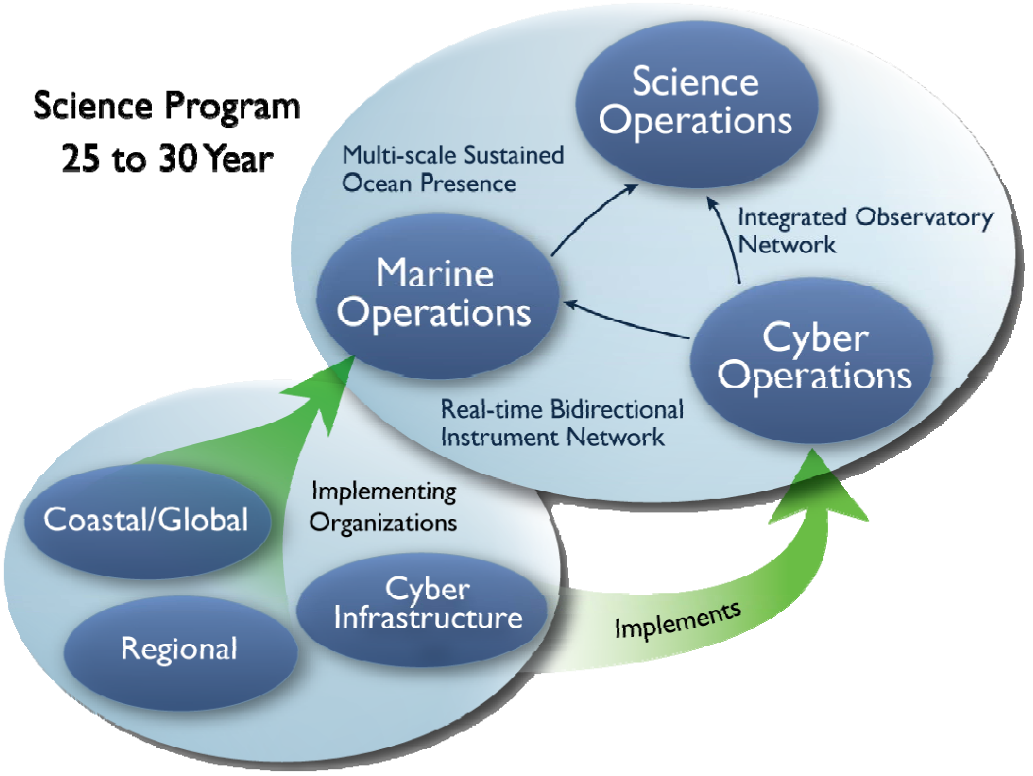

Benefits and Concerns Now

- Benefits
  - Environment per user (group)
  - On-demand access
  - “We don’t want to run datacenters!”
  - Capital expense -> operational expense
  - Growth and cost management
- Concerns
  - Performance: “Cloud computing offerings are good, but they do not cater to scientific needs”
  - Price and stability
  - Privacy

Performance (Hardware)

- Challenges
  - Big I/O degradation
  - Small CPU degradation
- Ethernet vs Infiniband
  - No OS bypass drivers, no infiniband
- New development
  - OS bypass drivers

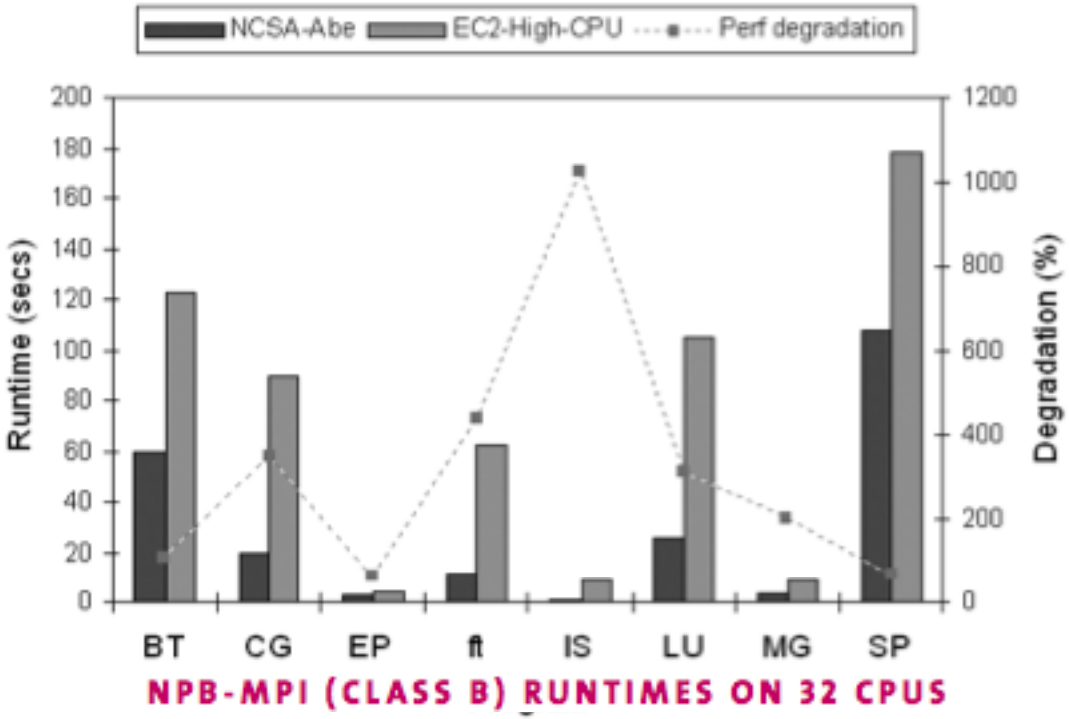

From *Walker, ;login: 2008*

Release: Xen 4.0 – UPDATED

Xen

Performance (Configuration)

- Trade-off: CPU vs. I/O
- VMM configuration
  - Sharing between VMs
  - “VMM latency”
  - Performance instability
- Multi-core opportunity
- A price performance trade-off

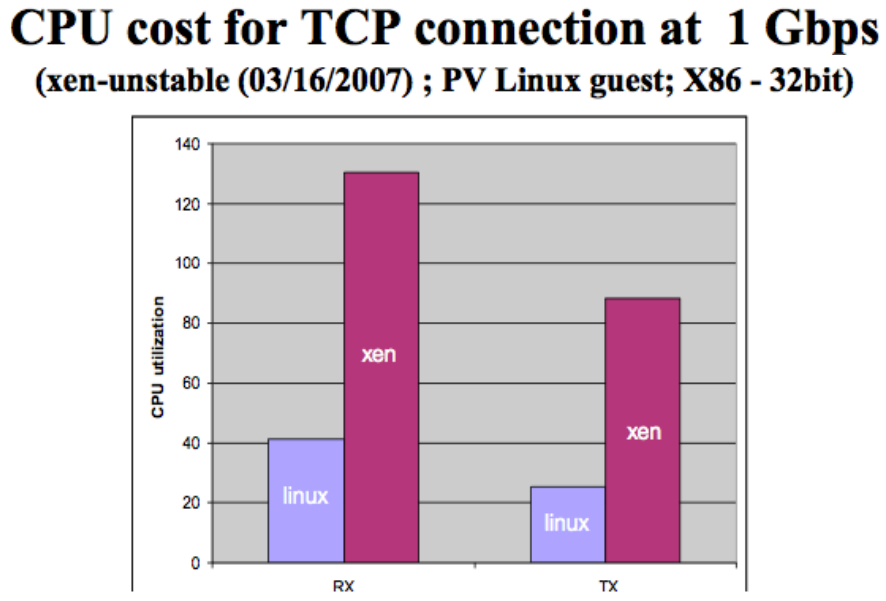

From *Santos et al., Xen Summit 2007*

Ultimately a change in mindset: “performance matters”!

From Performance...
...to Price-Performance

- “Instance” definitions for science
  - I/O oriented, well-described
- Co-location of instances
- To stack or not to stack
- Availability @ price point
  - CPU: on-demand, reserved, spot pricing
  - Data: high/low availability?
- Data access performance and availability
- Pricing
  - Finer and coarser grained

Availability, Utilization, and Cost/Price

- Most of science today is done in batch
- The cost of on-demand
  - Overprovisioning or request failure?
- Clouds + HTC = marriage made in heaven?
- Spot pricing

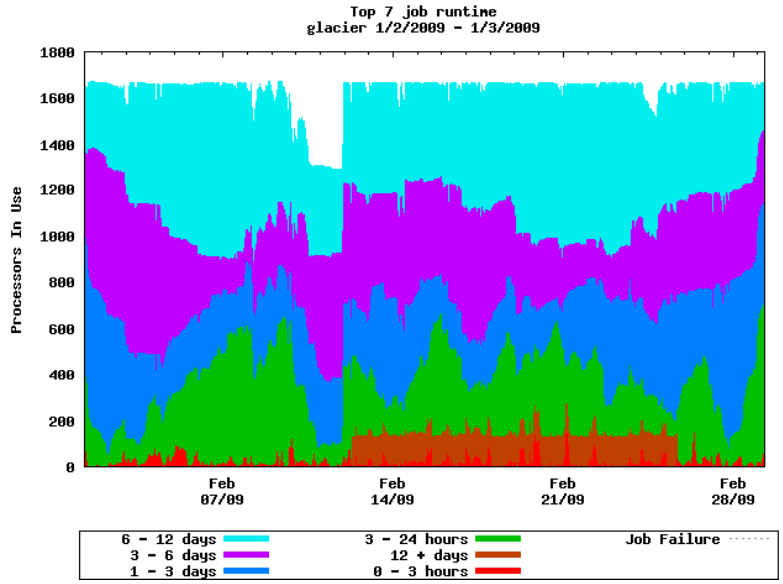

courtesy of Rob Simmonds, example of WestGrid utilization

Data in the Cloud

- Storage clouds and SANs
  - AWS Simple Storage Service (S3)
  - AWS Elastic Block Store (EBS)
- Challenges
  - Bandwidth performance and sharing
  - Sharing data between users
  - Sharing storage between instances
  - Availability
- Data Privacy (the really hard issue)

Descher et al., Retaining Data Control in Infrastructure Clouds, ARES (the International Dependability Conference), 2009.

Cloud Markets

- Is computing fungible?
- Can it be fungible?
  - Diverse paradigms: IaaS, PaaS, SaaS, and other aaS..
  - Interoperability
  - Comparison basis

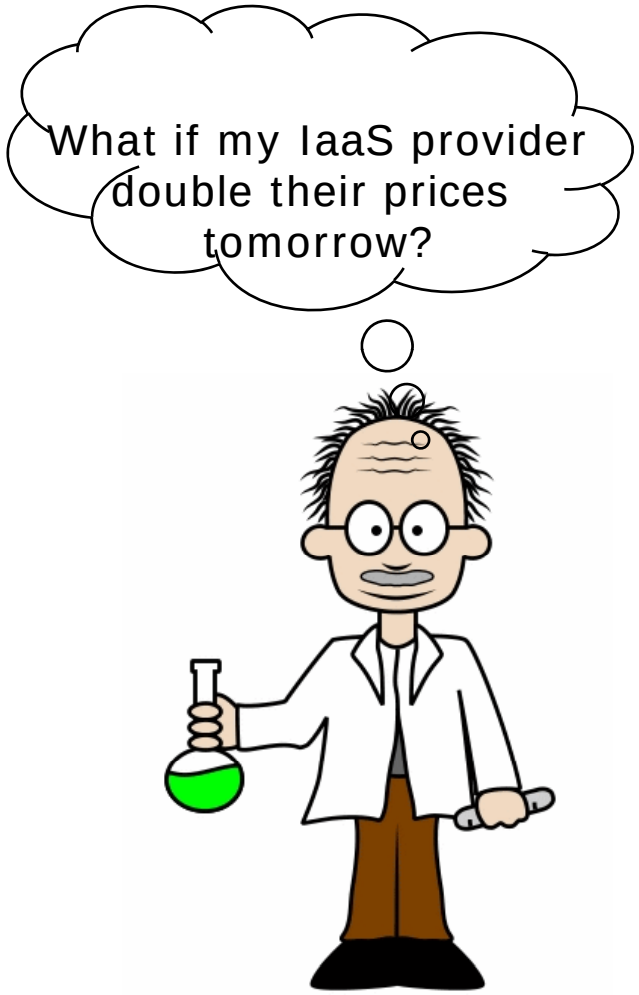

IaaS Cloud Interoperability

- Cloud standards
  - OCCI (OGF), OVF (DMTF), and many more...
  - Cloud-standards.org
- Cloud abstractions
  - Deltacloud, jcloud, libcloud, and many more...
- Appliances, not images
  - rBuilder, BCFG2, CohesiveFT, Puppet, and many more...

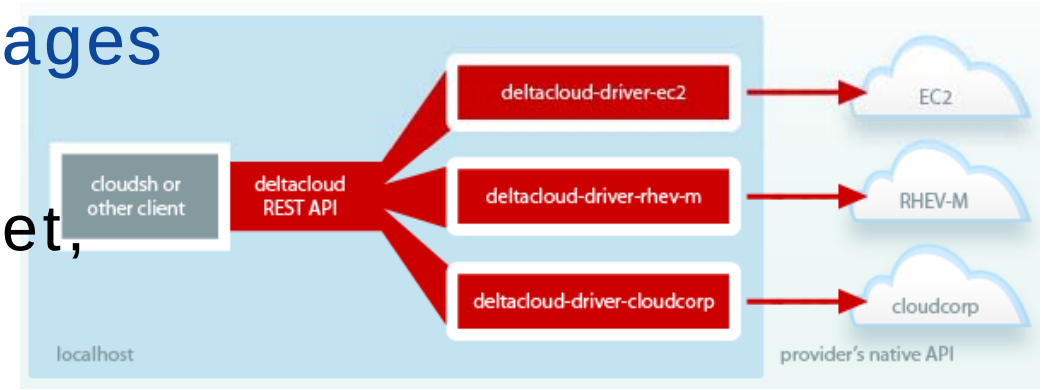

Can you give us a better deal?

- In terms of…
  - Performance, deployment, data access, price
- Based on relevant scientific benchmarks
- Comprehensive, current, and public

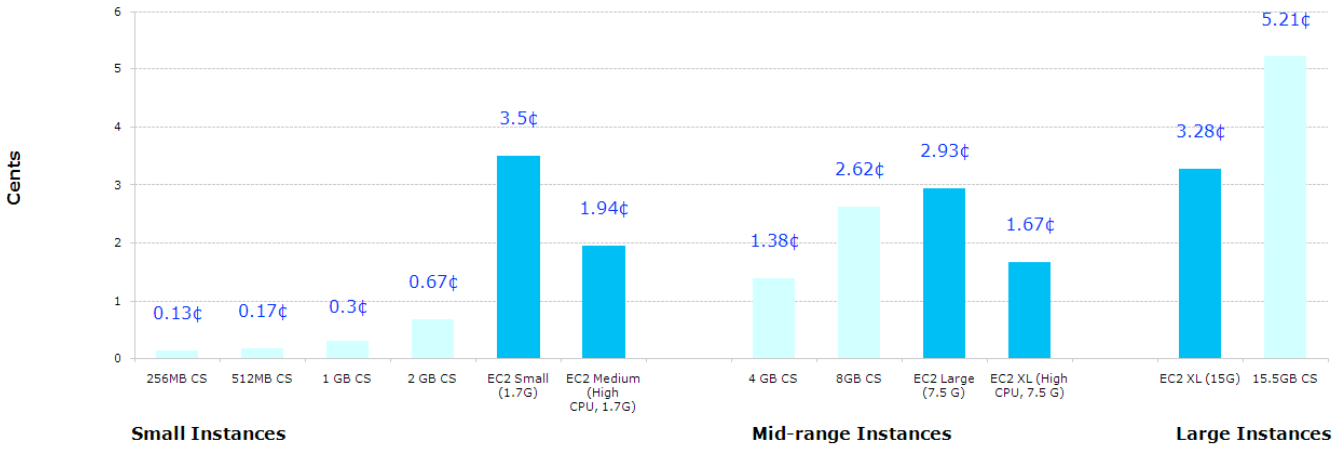

The Bitsource: CloudServers vs EC2 LKC Cost by Instance

Science Cloud Ecosystem

- Goal: time to science in clouds -> zero
- Scientific appliances
- New tools
  - “turnkey clusters”, cloud bursting, etc.
  - Open source important
  - Change in the mindset
- Education and adaptation
  - “Profiles”
  - New paradigm requires new approaches
  - Teach old dogs some new tricks

The Magellan Project

- Joint project: ALCF and NERSC
- Funded by DOE/ASCR ARRA

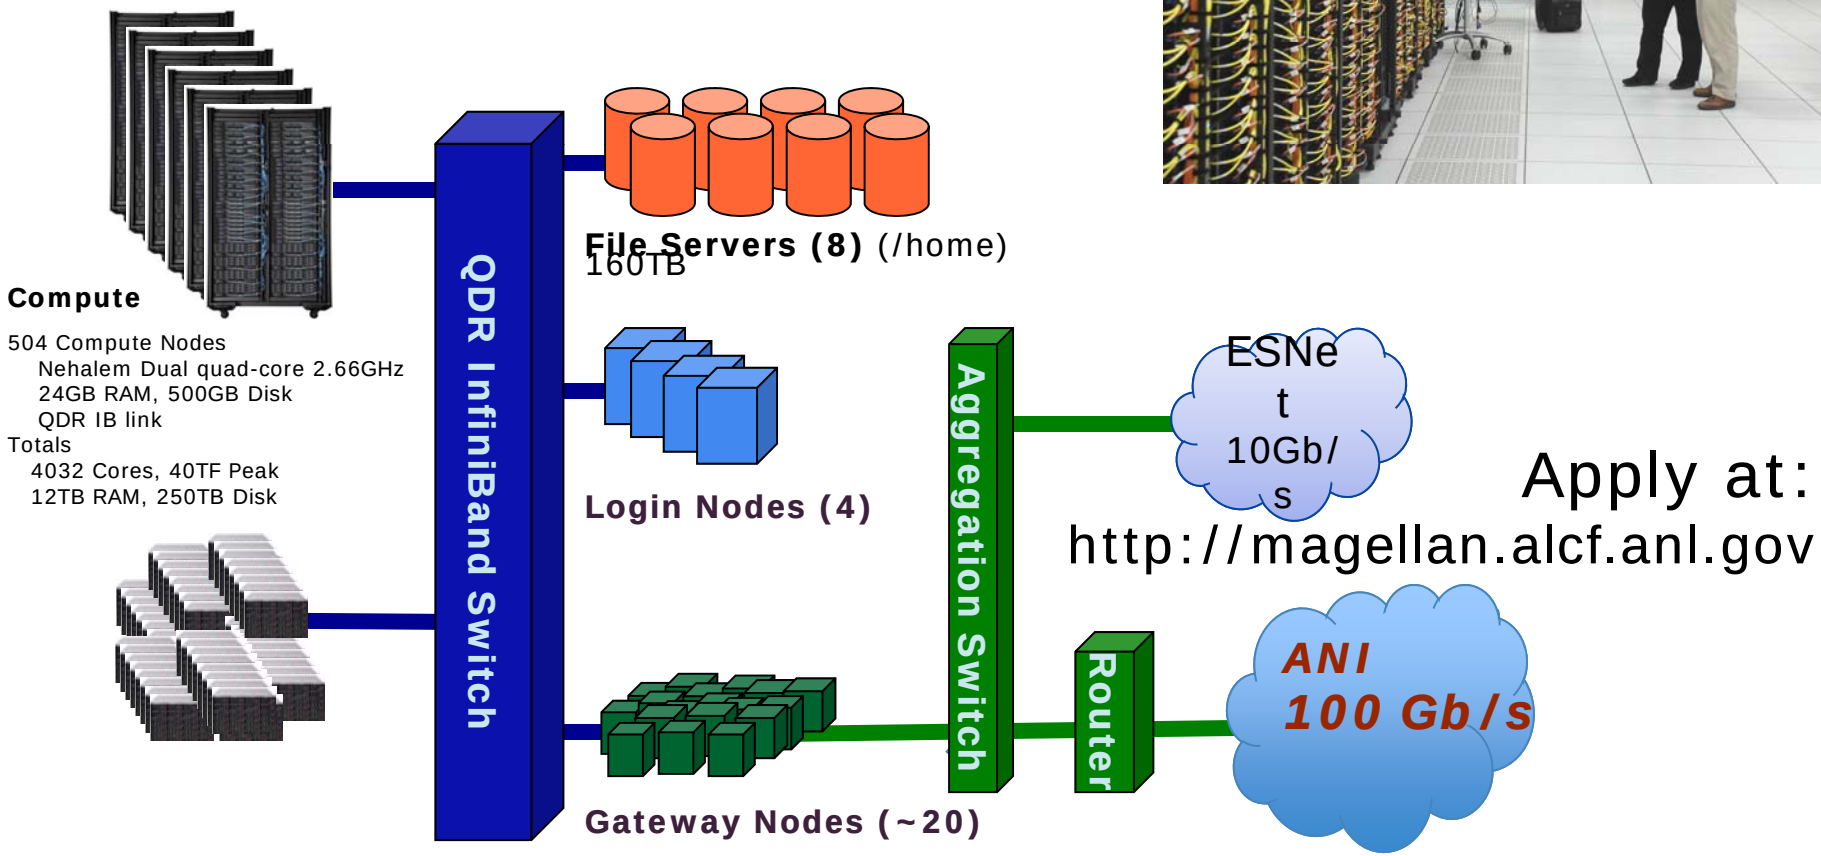

Apply at:
http://magellan.alcf.anl.gov

The FutureGrid Project

- NSF-funded experimental testbed (incl. clouds)
- ~6000 total cores connected by private network

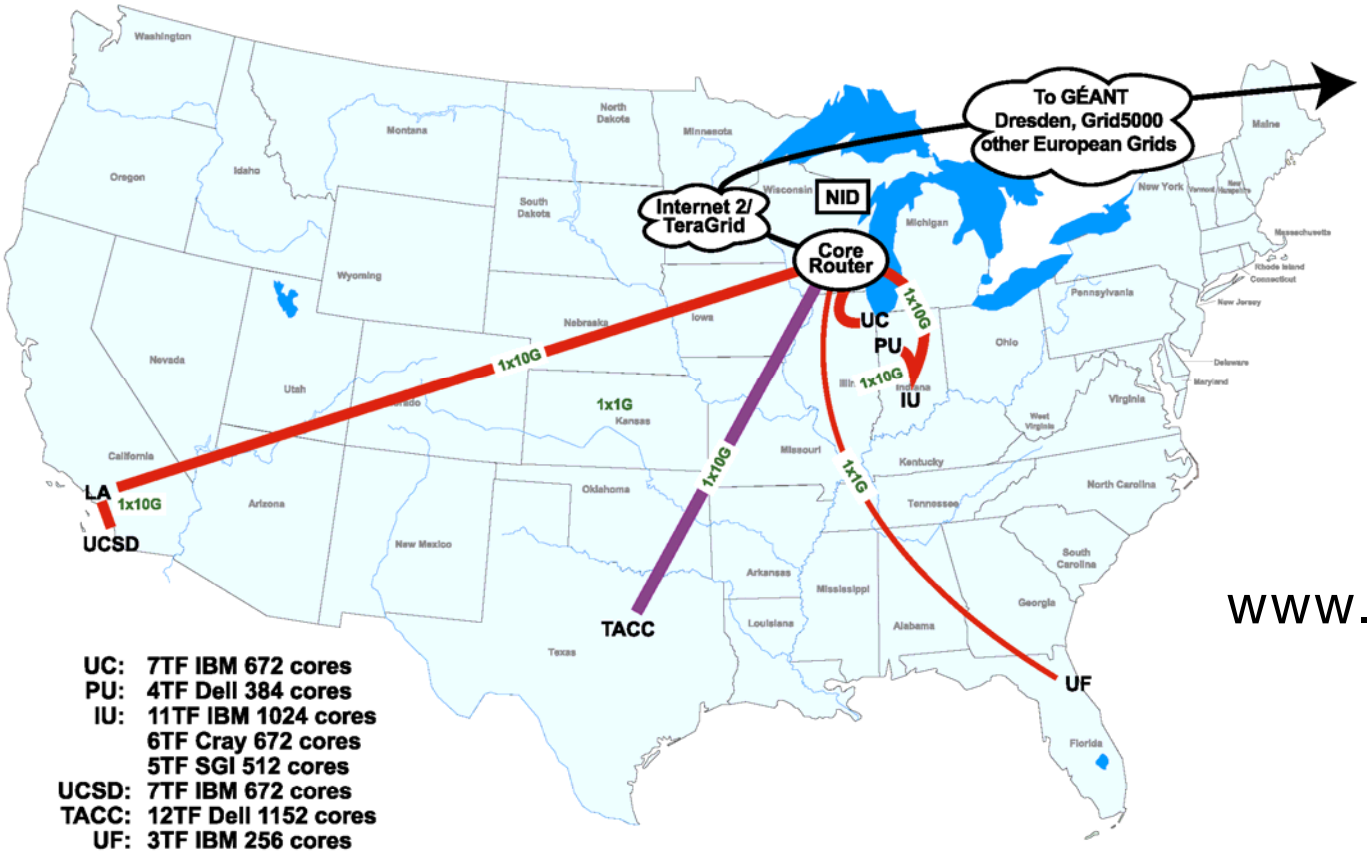

UC: 7TF IBM 672 cores
PU: 4TF Dell 384 cores
IU: 11TF IBM 1024 cores
6TF Cray 672 cores
5TF SGI 512 cores
UCSD: 7TF IBM 672 cores
TACC: 12TF Dell 1152 cores
UF: 3TF IBM 256 cores

Apply at:
www.futuregrid.org

Parting Thoughts

- Opinions split into extremes:
  - “Cloud computing is done!”
  - “Infrastructure for toy problems”
- The truth is in the middle
  - We know that IaaS represents a viable paradigm for a set of scientific applications…
  - …it will take more work to enlarge that set
  - Experimentation and open source are a catalyst
- Challenge: let’s make it work for science!