Fermilab

Managed by Fermi Research Alliance, LLC for the U.S. Department of Energy Office of Science

Fermilab HEPCloud Facility: Data-Intensive Computing In the Cloud

Steven C. Timm

OpenNebulaConf 2015

21 October 2015

Presentation Outline

- History of FermiCloud Project and OpenNebula
- Motivation for expanding to commercial clouds
- Goals of the HEPCloud project
- Use Cases and Architecture
- Significant Developments thus far

FermiCloud and OpenNebula

- FermiCloud project began in 2010
  - At that time data costs to commercial clouds were considered prohibitive—so focus on private cloud.
- First OpenNebula installed in November of 2010
- Quickly grew to be our only cloud system.
- X.509 authentication modules for OpenNebula >=3
- Currently running OpenNebula 4.8
- We have 29 bare metal hosts for interactive/special purpose VM's and 280 older 8-core hosts for running batch virtual machines.—maintained by ~0.2 of an employee (Me).
- OpenNebula developers have been very cooperative throughout, always willing to take our patches and features.
  - Key reason why we chose and stayed with OpenNebula

Important outcomes of FermiCloud Project

- MPI in VM's across virtualized Infiniband
- Deployment of X.509 authentication and authorization
- Secrets repository for virtual machines.
- Extensive study of distributed file system performance in VMs
- Federated cloud using Amazon and OpenNebula VM's
- Idle virtual machine detection software development
- Testbed for 100GBit WAN and IPv6 efforts
- Developed automated imaging system to serve OpenNebula and AWS
- Important testing and integration facility for the Open Science Grid and also our own middleware developers at Fermilab.

Why expand to commercial clouds?

- Best and most efficient internal IaaS still doesn't solve key problems for US R+D agencies:
- Still have to have a computer building, people to manage the computers, and procure computers on long government procurement life-cycle
  - US funding agencies reluctant to build new computing centers
  - Procurement/deployment long and getting longer
- Demand shows a peaking pattern
  - We see steady DC usage from our users plus temporary peaks of 4-5 times the demand
  - One current use case requires 56000 cores for a month, 4x what we have in our facility. Can't buy machines or build buildings that fast.
  - Expect that opportunistic cycles on Grid will be less available going forward as well—although that is NOT true yet.
- The users we serve want to do it and need the cycles.

What keeps us up at night! (II)

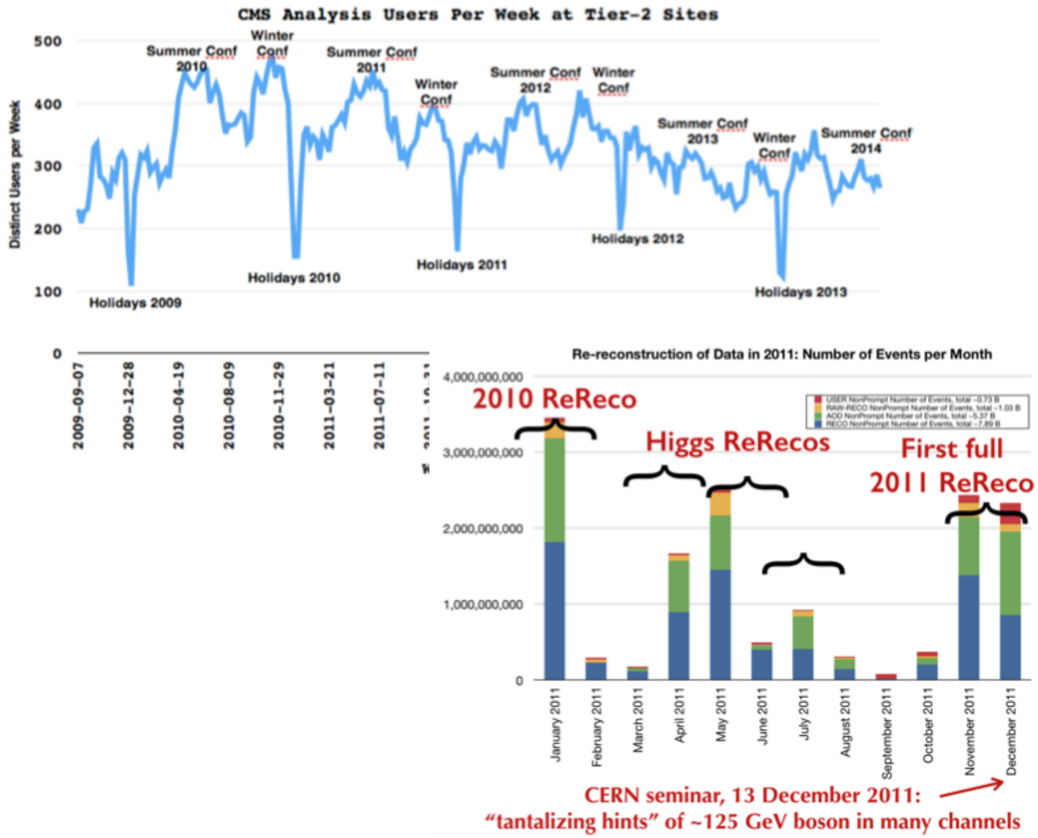

- Activity of experiments is not constant
  - It varies significantly with external triggers
    - Operation schedule
    - Conference schedule
    - Holidays, vacation time, etc.
- Question: How can we provision resources efficiently? ➜ **ELASTICITY**

Fermilab

3 Gabriele Garzoglio, Oliver Gutsche I CHEP2015: Diversity in Computing Technologies and Strategies for Dynamic Resource Allocation 15. April 2015

The Charge of the Fermilab HEPCloud Facility Project

- Extend the current facility services to transparently run on a mix of disparate resources including
  - Commercial and community clouds
  - Grid Federations
  - HPC Centers
- Perform FULL SPECTRUM of computing including
  - Data-intensive simulation and reconstruction
  - At production scale whether local, remote, or both.
- Allow facility to respond to peaks in demand without overprovisioning local resources
- Make all of this transparent to users

Use Case I: NOvA experiment at Fermilab

- Physics Goals:
  - New: Run the most data-intensive NOvA applications on AWS
  - Event Reconstruction—six campaigns
    - 10000 input files: Each 250MB input, 250 MB output, 3hrs/job
  - Beam Simulation with Flux Files—ten campaigns
    - 38000 input files, 5 hrs/job
  - Will add 3-4 other use cases from other experiments
- Sustainability goals
  - Total expected AWS usage 2.1 million hours (100x 2014 test)
  - For scale, NOvA alone ran 10.2 million hours in 2014
  - 145 million hours/year available on non-CMS FermiGrid.
- Of interest: Final stage of NOvA analysis requires 256GB of RAM (maybe more).. Prototyped this on AWS too.

Use Case II: CMS Experiment @ CERN

- New high-energy run @ LHC started earlier in 2015
- New energy, new beam spot, have to re-do lots of simulation
- Goal—1 Billion events of simulation in one month
  - Would take ~4 months to run this at Fermilab US Tier 1.
  - Estimate 56000 CPU cores DC for one month
  - 800 TB of output data generated.
- Four steps:
  - Generation—throw dice to generate desired event
  - Simulation—how will it decay in the detector
  - Digitization—what will the signal look like in the detector
  - Reconstruction—take electronic signals and make tracks in 3D.
- May also add reconstruction of experimental data.

HEPCloud Facility Architecture

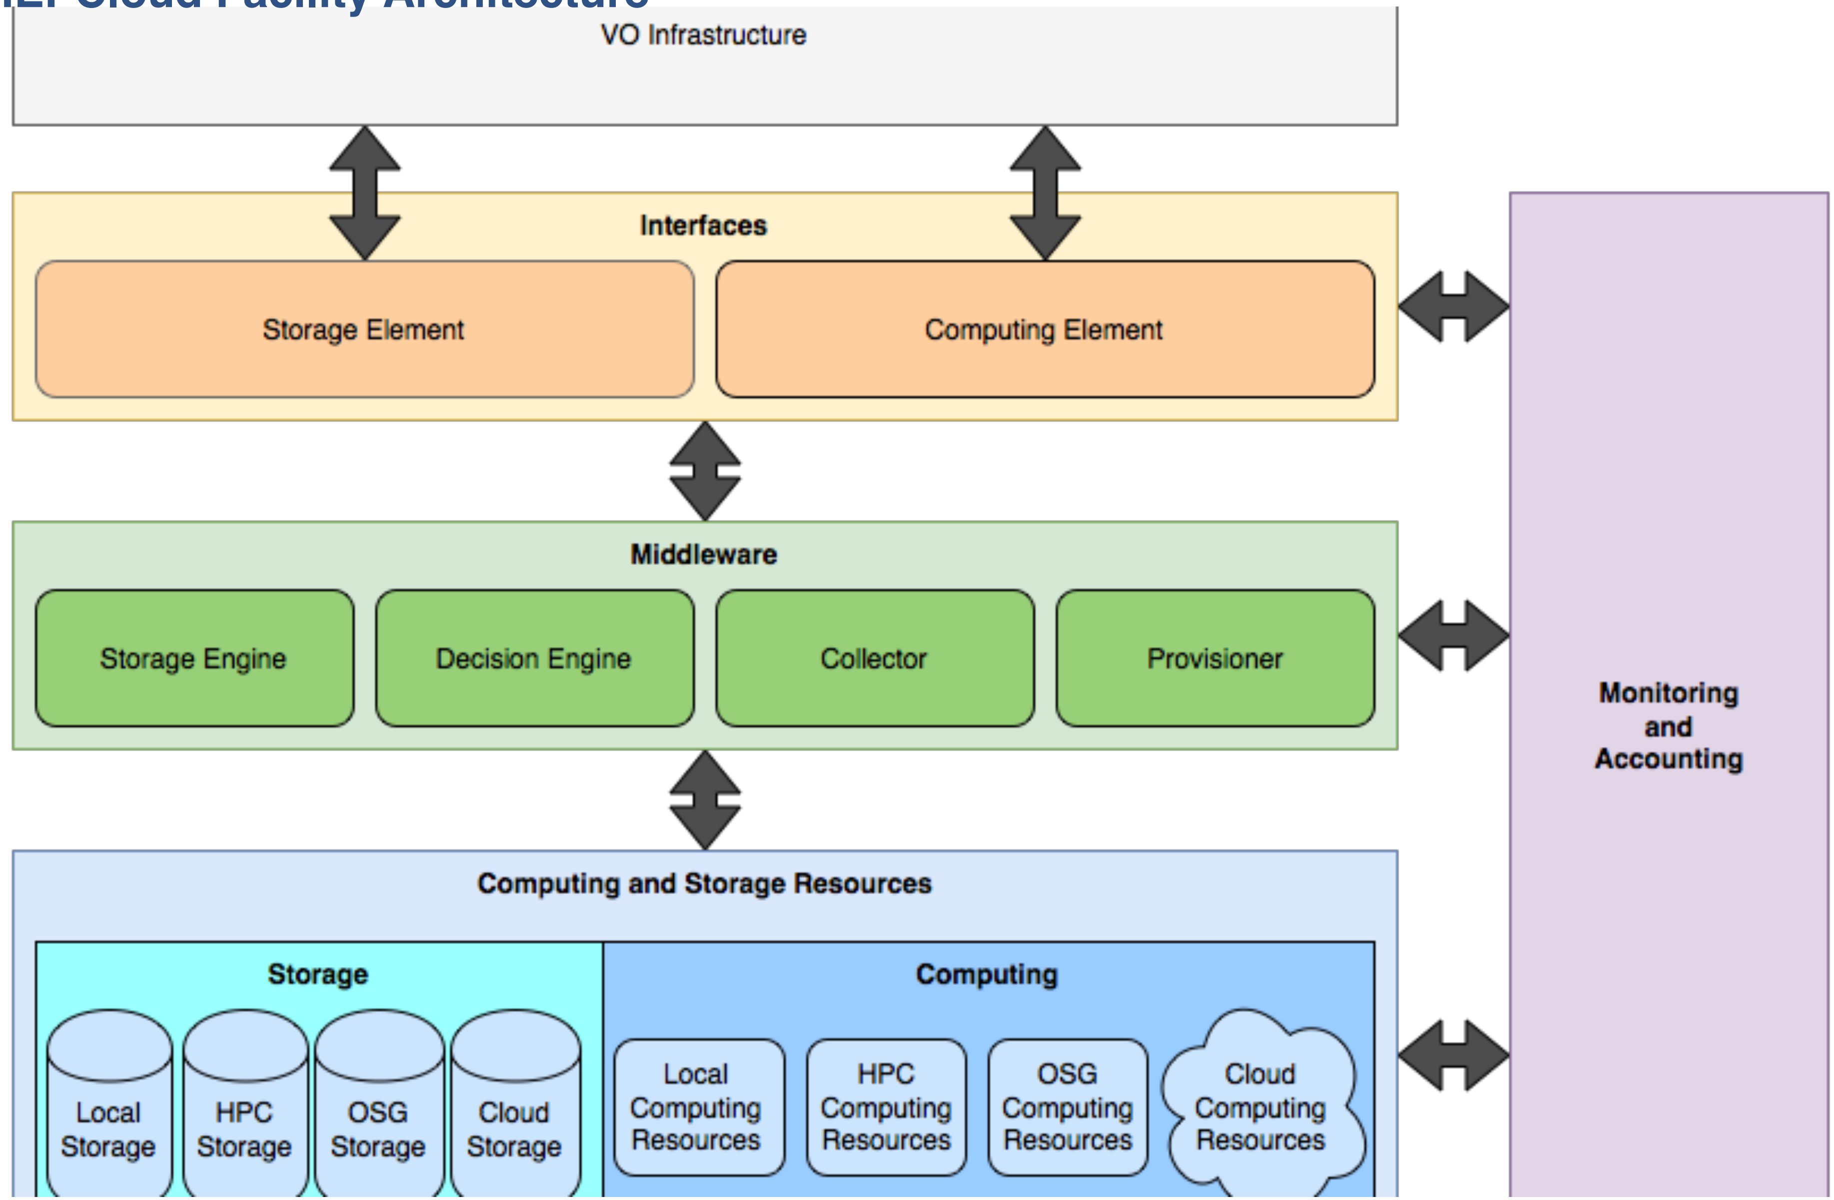

Network Topology

- Major US research networks (ESNet, Internet2) peer with cloud providers at nearby points of presence, up to 100Gbps in Western US.
- Amazon “Direct Connect” allows to extend corporate network directly into Amazon using BGP and VPC gateway
- Fermilab decided it was better to not do this
  - Instead use normal Internet routing and Amazon Internet gateway—firewall rules block non-Fermilab traffic.
- High latency requires that multiple streams are used for data transfers
- Have demonstrated up to 7Gbps sending files from AWS to Fermilab storage systems.
  - This is 2X what we need right now but think we can do better.

The Decision Engine

- Historically the glideinWMS provisioning asks:
  – Does this class of jobs need slots?
  – If yes try to get as many as needed on VO-by-VO basis.
- Now need smarter algorithm:

  Optimize job location based on cost and completion time
  – Burst to public cloud if:
    - The local Fermilab facility is full
    - And it's likely to stay full for the length of time it takes to get slots started on Amazon
    - And there are jobs that are technically suited to run on the cloud
    - And the group has budget and priority to run on the cloud
    - BUT avoid letting one graduate student burn all the cloud time.
- Also consider existence of "friendly" grid sites and HPC
- Major new intellectual capital of the project is here.

Prediction Engine 1

- Have to plan workflows more carefully when spending $$
- How much CPU time (relative to bare metal machines)
- How much stage in / stage out time?
- Cheaper to stage to cloud storage or to read/write across WAN?
- Most importantly—what spot price to bid for resource?
- Resources available for as low as 12% of “on demand” price
- If someone bids higher than you and there are no other available resources, you get pre-empted
- Much in the literature on how to do this BUT
  - They assume checkpointable application, ours are not.
  - They assume short application, ours are 8-24 hours.
- Christmas is coming soon—high demand period.

Prediction Engine 2

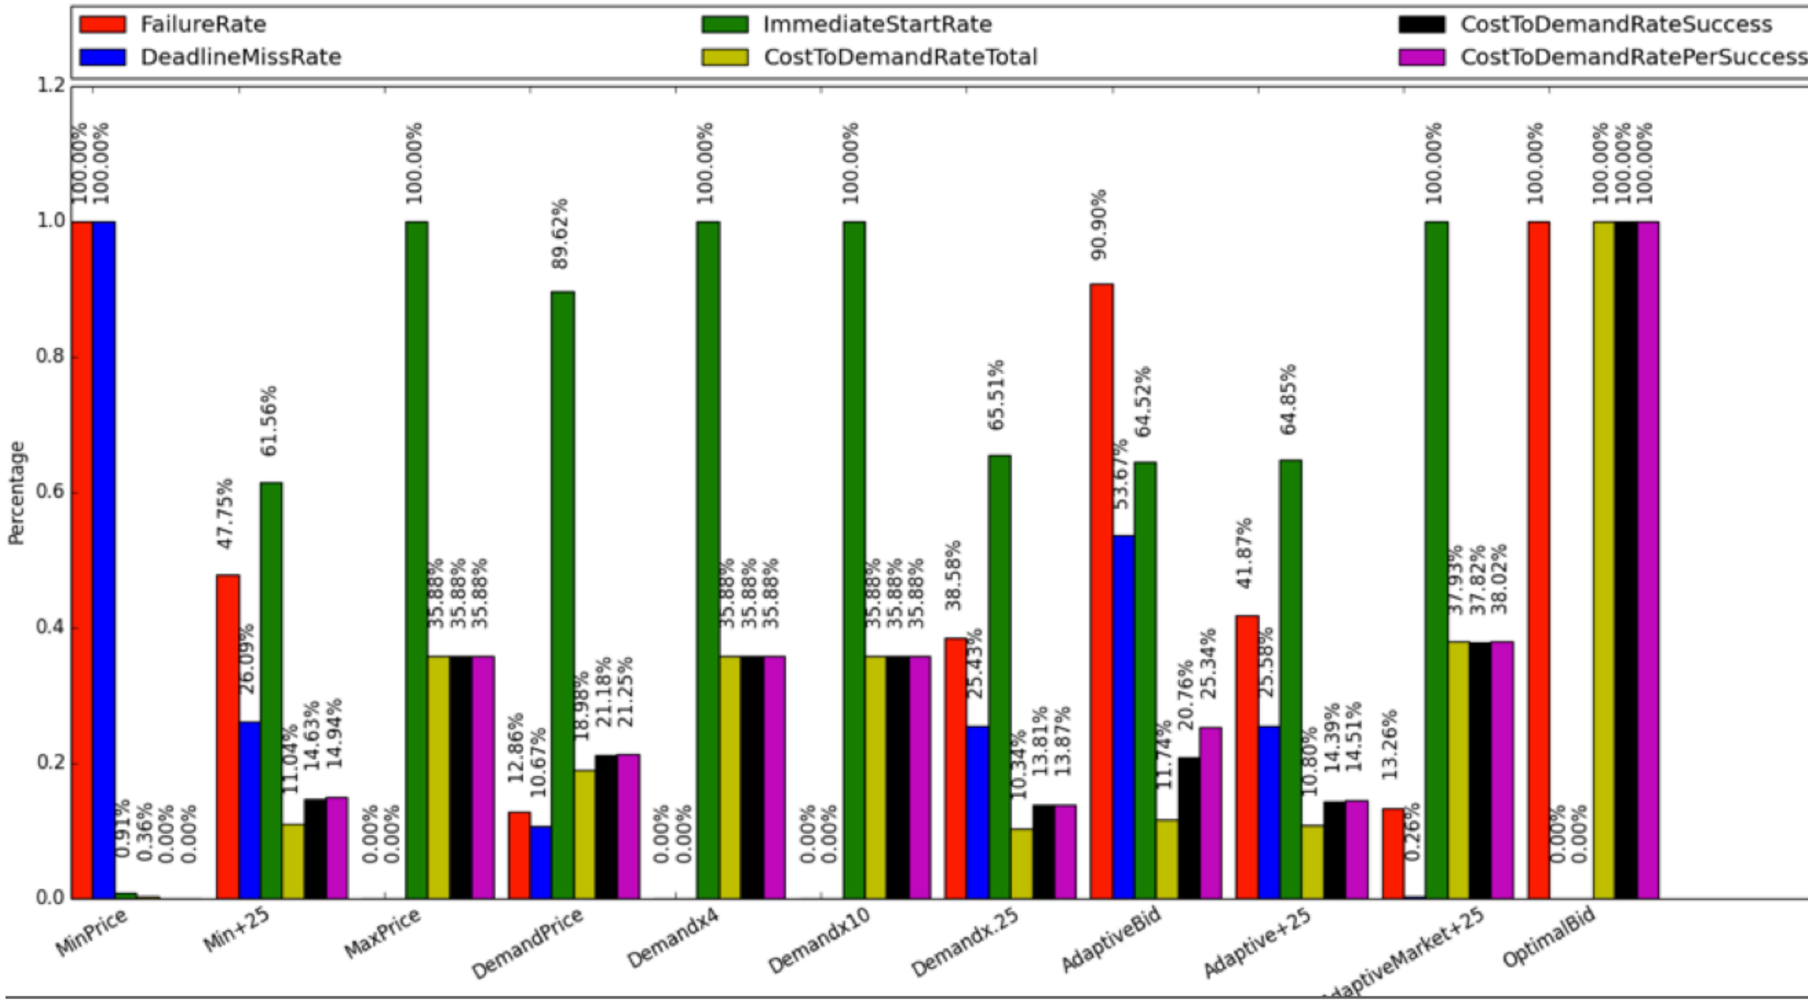

On-Demand Services

- Need to transport software and data to the Cloud
- Auto-scalable squid servers, deploying and destroying in 30 seconds using CloudFormation script
- Squid—basically a stateless service, easy to scale up and down.
- Route53—makes a DNS alias to tell us where to find it.
- Scale tested—good for 3 Gb/s, 500000 requests/minute

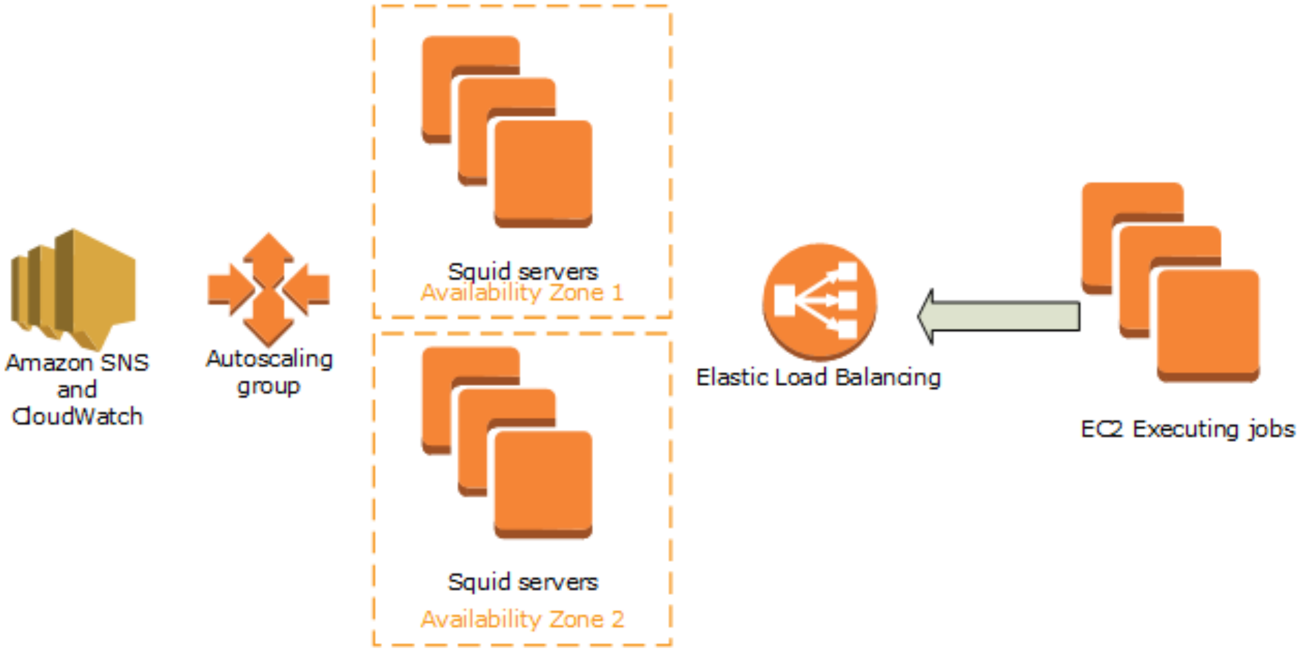

Fermilab

Future—Cloud Computing in US Science Labs

- New roles for Grid collaborations such as Open Science Grid in USA
  - Aggregate access to commercial + community clouds
  - Aggregate access to government supercomputing centers
- Economies of scale are hard to beat
  - Even for a site like Fermilab with huge network pipe, big computing center already built, and lots of access to cheap power.
  - Hard to do the real calculation in a way that everyone agrees
- Many collaborations including ours will have Cloud in the name but include significant non-cloud activities @HPC.
  - Big opportunity for those who really needed HPC anyway
  - Big resource for high-throughput computing

Future: Cloud Computing at US Science Labs

- Nevertheless there is still space for open-source IaaS at local facilities.
  - Important to keep a big chunk of computing close to storage
  - Flexibility in server provisioning and reprovisioning
  - R+D for the big cloud efforts
  - Space for applications that don't fit standard HTC model.
    - High memory, high disk, high IOPS
- CERN going all-in with IaaS cloud is an anomaly,
  - Most US facilities won't do that.
- In US—Amazon dominates for now because they have done homework of how to sell to government (incl. security).
- Funding agencies will require diversity and open standards.

Summary:

- Good cloud projects start with good policy, good security stance, and good networks
- Flexible provisioning systems are key to expanding from local clusters to grid to cloud to HPC.
- The longer you have a cloud system the more people think of new ways to use it and do computing they could not otherwise have done.
- HEPCloud has successful tests of both of its use cases and is ramping up major scale tests now.
- High Energy Physics demand for computing cycles is potentially near-infinite. We need to make use of all possible platforms.

Muchas Gracias / Moltes Gràcies / Thank You

amazon web services™

Scientific Linux

KiSTi www.kisti.re.kr

ILLINOIS INSTITUTE OF TECHNOLOGY

OpenNebula

CMS

NOvA

HTCondor High Throughput Computing

Fermilab