The HEP Cloud Facility: elastic computing for High Energy Physics

Gabriele Garzoglio, for the Fermilab HEPCloud Team
HEP Cloud Facility co-Project Manager

CHEP 2016
Oct 11, 2016

Drivers of Evolution: Capacity / Cost / Elasticity

HEP needs: 10-100 x today capacity

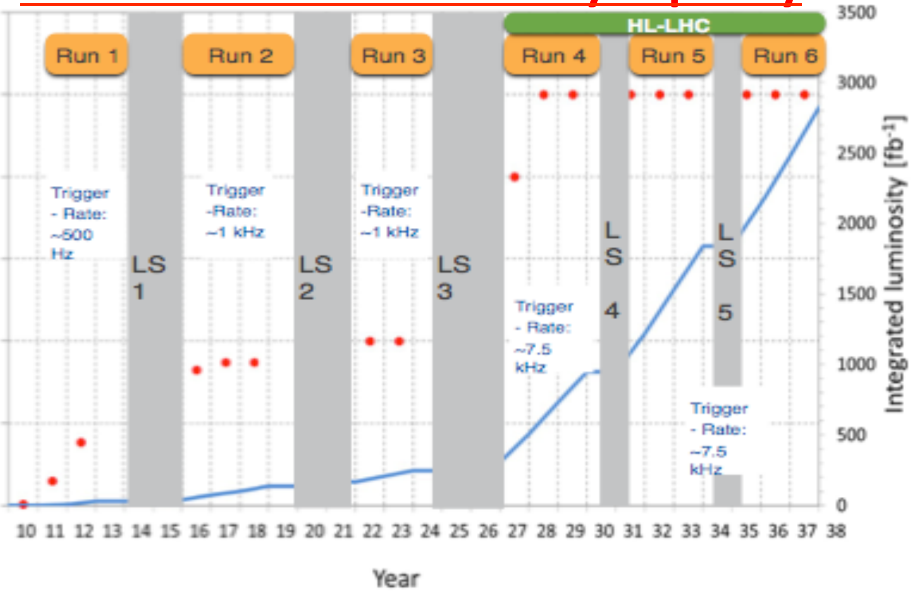

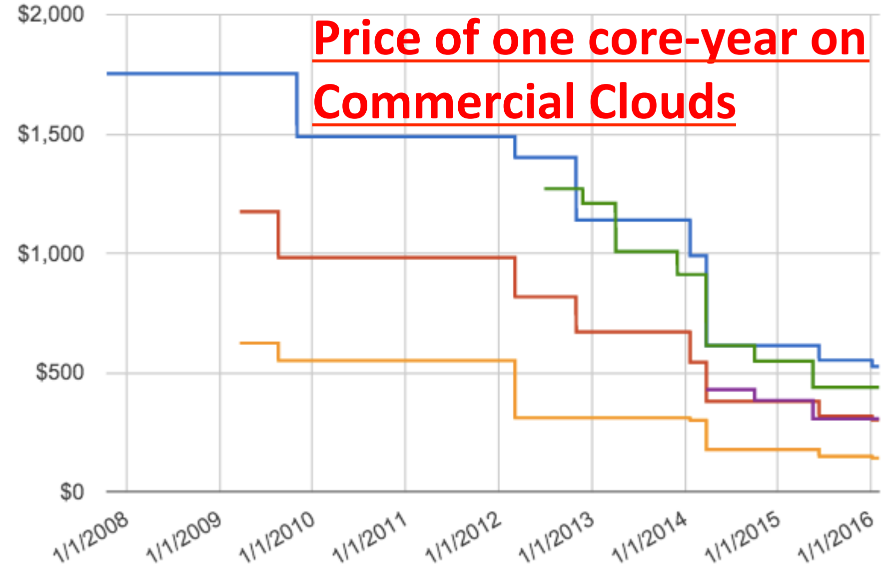

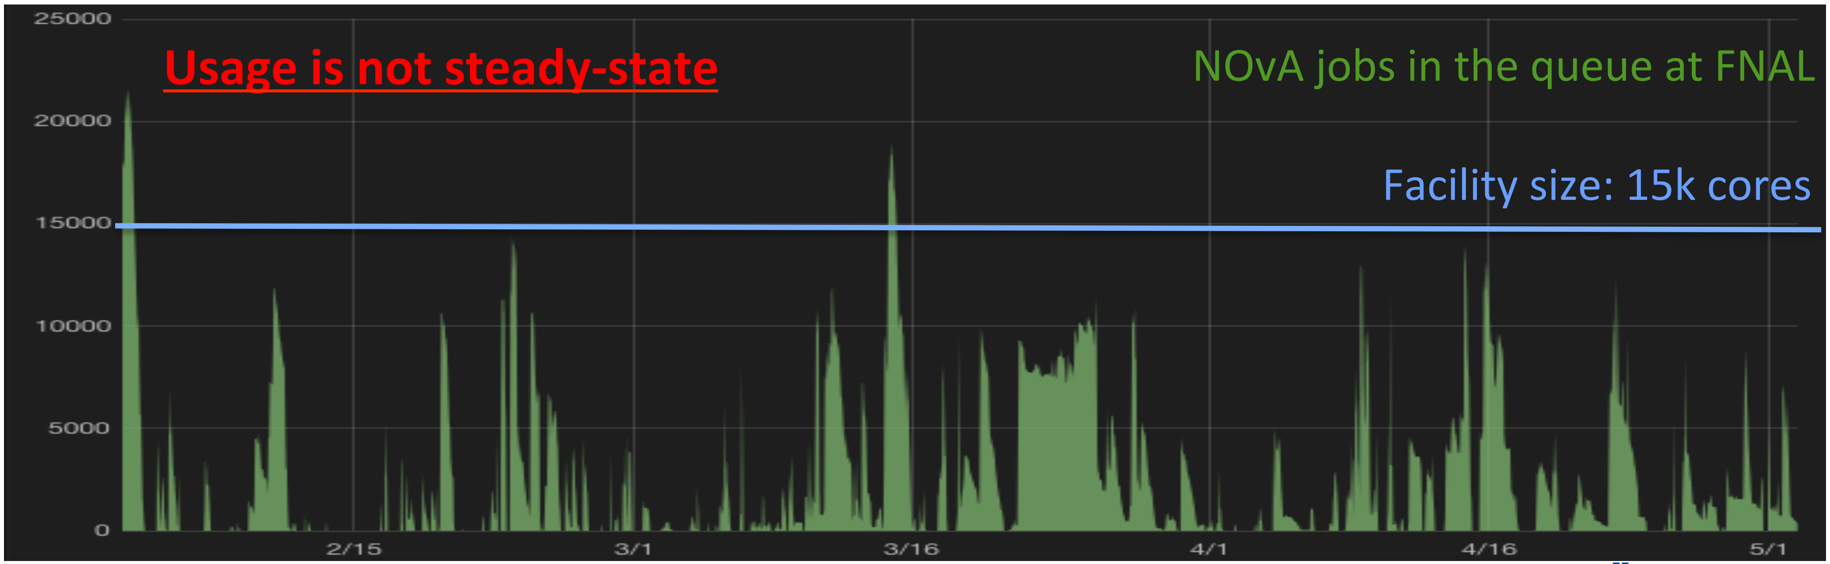

Vision for Facility Evolution

- Strategic Plan for U.S. Particle Physics (P5 Report)

  *Rapidly evolving computer architectures and **increasing data volumes require** effective crosscutting **solutions** that are being **developed** in other science disciplines and **in industry**.*

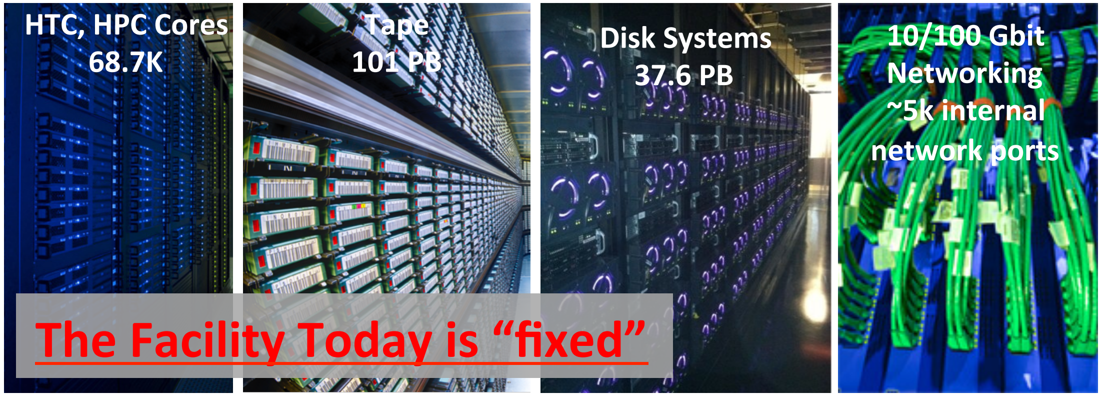

- HEP Cloud Vision Statement
  - **HEPCloud** is envisioned as a **portal** to an ecosystem of **diverse computing resources** commercial or academic
  - Provides "**complete solutions**" to users, with agreed upon levels of service
  - The Facility routes to **local or remote resources** based on workflow requirements, cost, and efficiency of accessing various resources
  - Manages allocations of users to target compute engines

- Pilot project to explore feasibility, capability of HEPCloud
  - Goal of moving into **production during FY18**
  - Seed money provided by industry

HEP Cloud Architecture

Overview

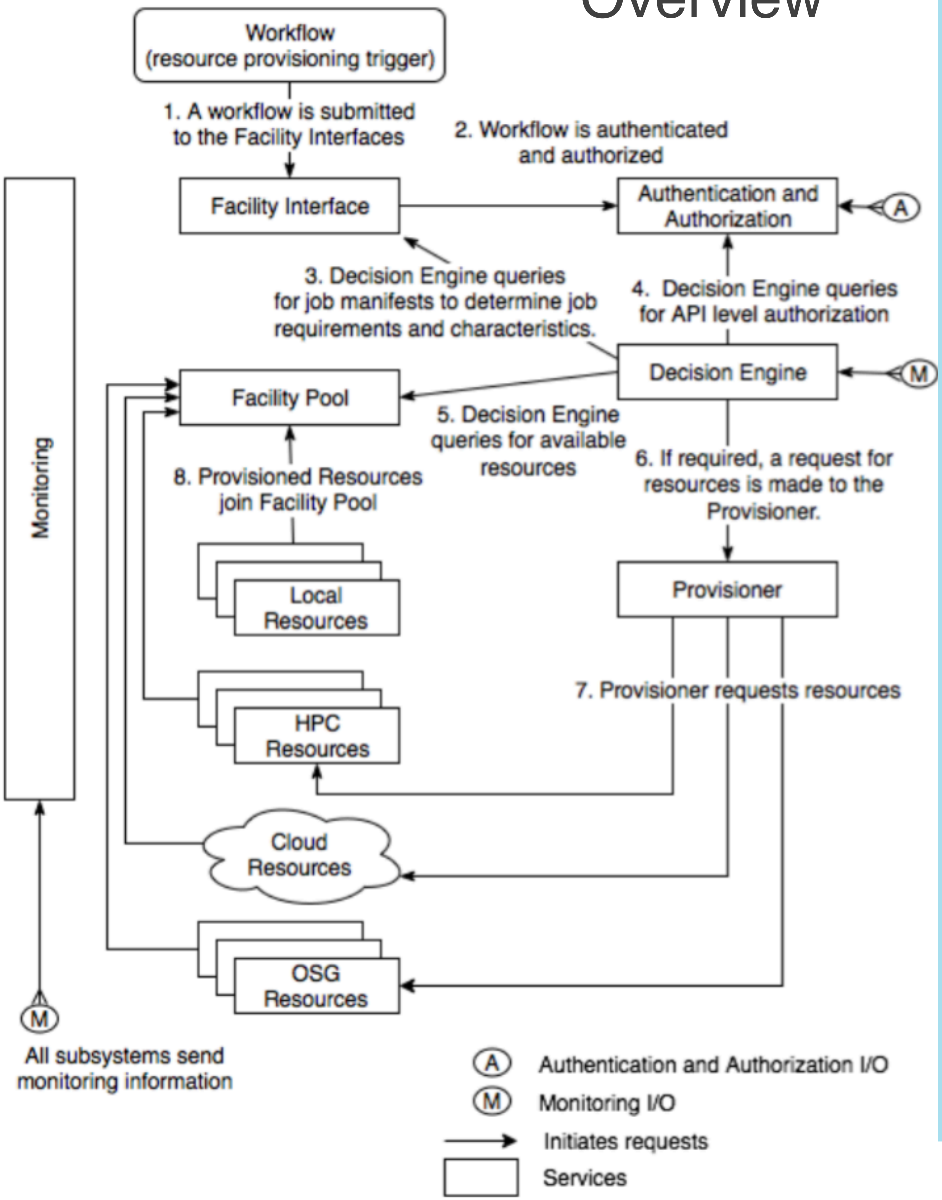

External Relationships

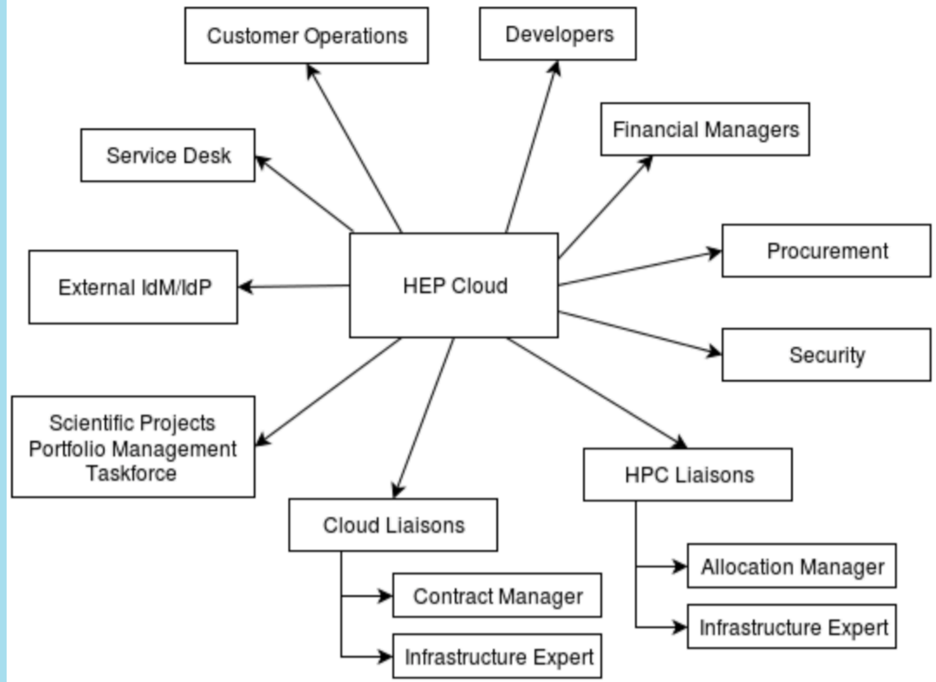

HEP Cloud Architecture

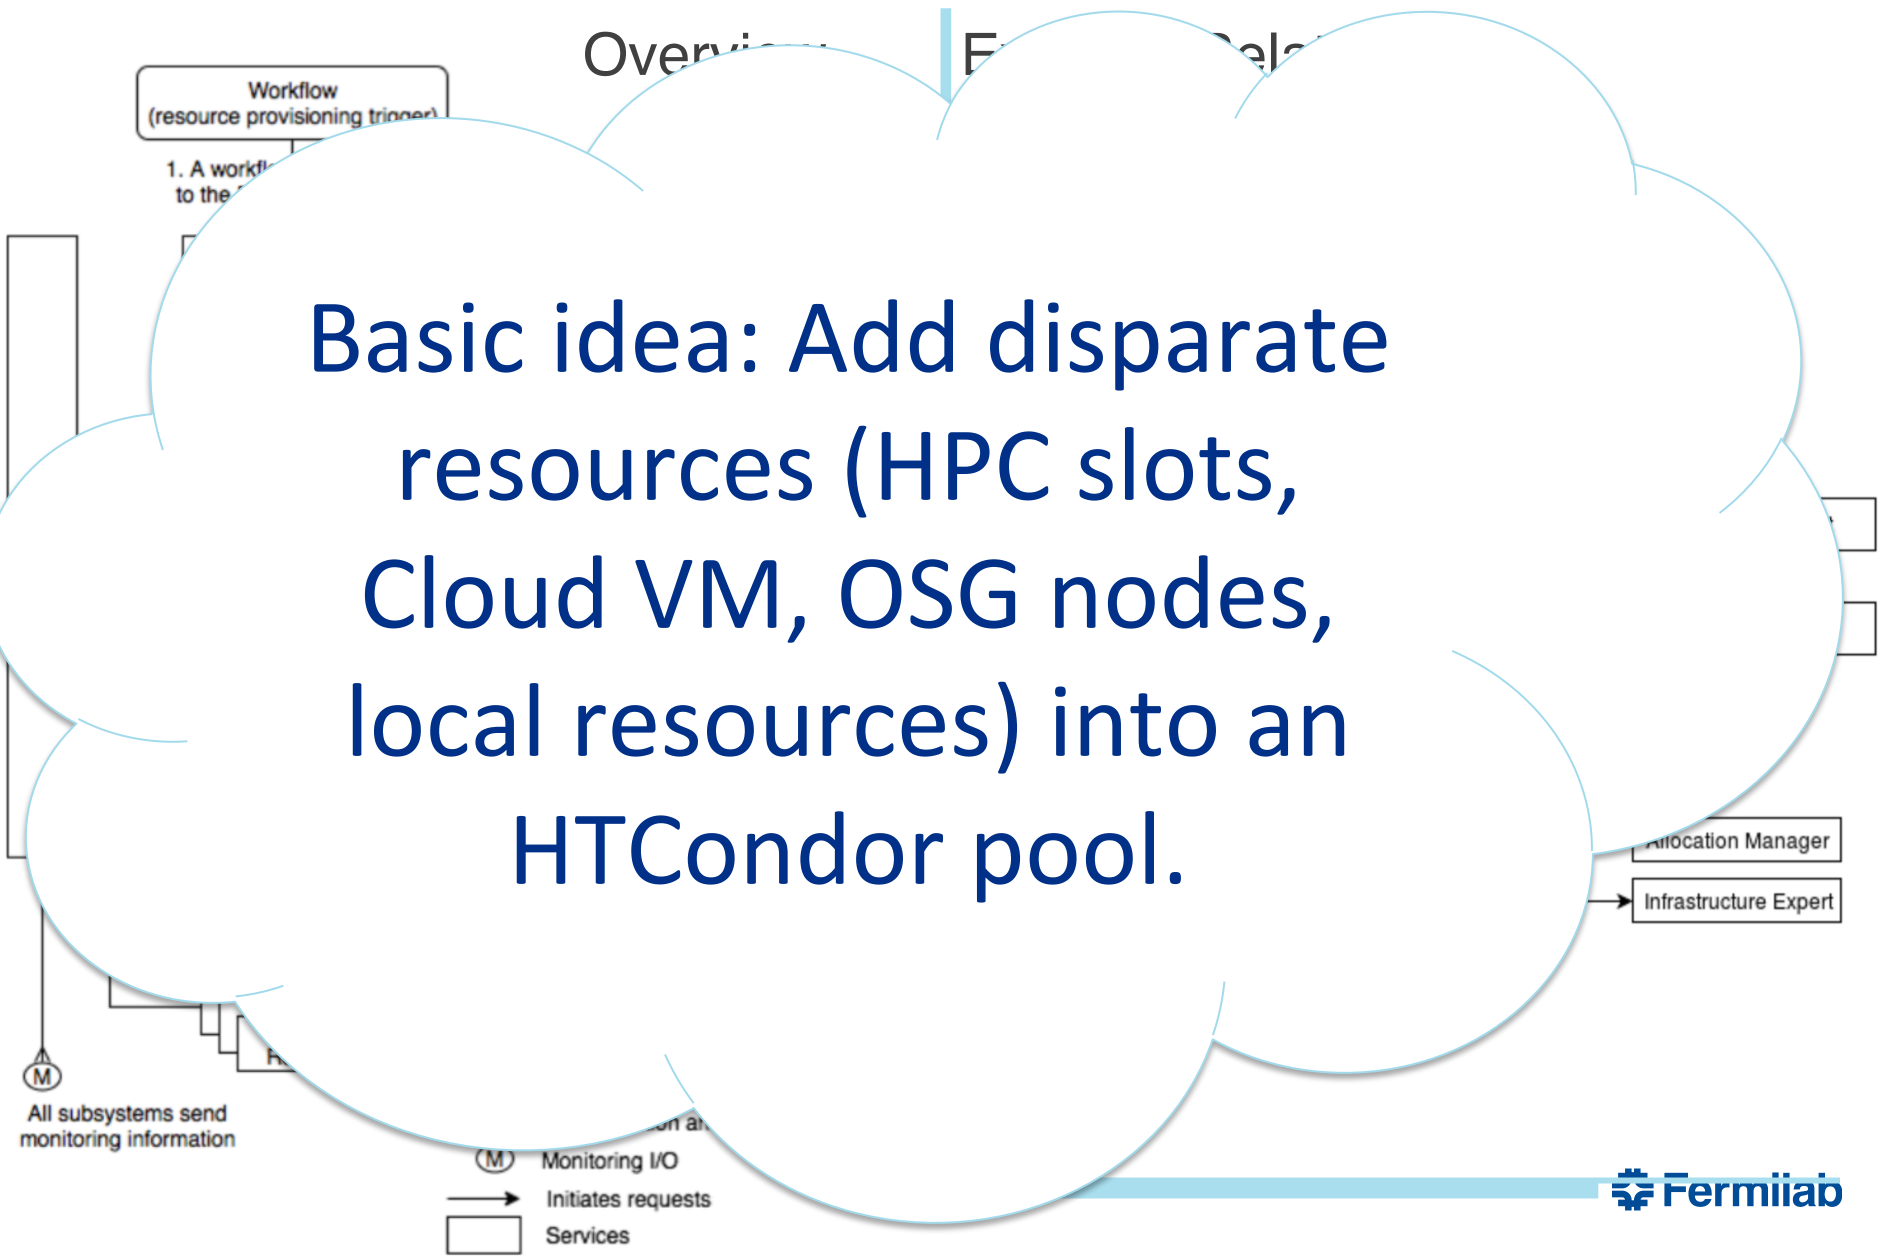

Basic idea: Add disparate resources (HPC slots, Cloud VM, OSG nodes, local resources) into an HTCondor pool.

Classes of External Resource Providers

Grid

- Virtual Organizations (VOs) of users trusted by Grid sites
- VOs get allocations ➜ **Pledges**
  - –Unused allocations: opportunistic resources

"Things you borrow"

Trust Federation

Cloud

- Community Clouds - Similar trust federation to Grids
- Commercial Clouds - **Pay-As-You-Go** model
  - Strongly accounted
  - Near-infinite capacity ➜ **Elasticity**
  - Spot price market

"Things you rent"

Economic Model

HPC

- Researchers granted access to HPC installations
- Peer review committees award **Allocations**
  - Awards model designed for individual PIs rather than large collaborations

"Things you are given"

Grant Allocation

Fermilab HEPCloud: expanding to HPC

- **Early steps**: adapt HTC workflows to HPC facilities
  - MicroBooNE production on Cori @ NERSC
    - Successfully downloaded the entire MicroBooNE release, including LArSoft and the art framework onto Cori, using Shifter from dockerhub.
    - **Executed single node tests of MicroBooNE Monte Carlo production**, reading from and writing to the global scratch file system through the container
  - Pythia on Mira @ ALCF: multi-parameter tuning of event generators using collider data
    - MPI + multi-threading to execute **32k instances** of Pythia and the Rivet analysis suite
    - Spirit of code-sharing – leveraged CMS contributions to multi-thread Pythia
  - CMS production: Provisioned resources and executed real GEN-SIM workflows at small scale
  - Plans for next year: mu2e, NOvA, MicroBooNE, CMS …

Fermilab HEPCloud: expanding to the Cloud

- Where to start?
  - Market leader: **Amazon Web Services (AWS)**
  - Now integrating **Google Compute Engine** (for Super Computing 16)

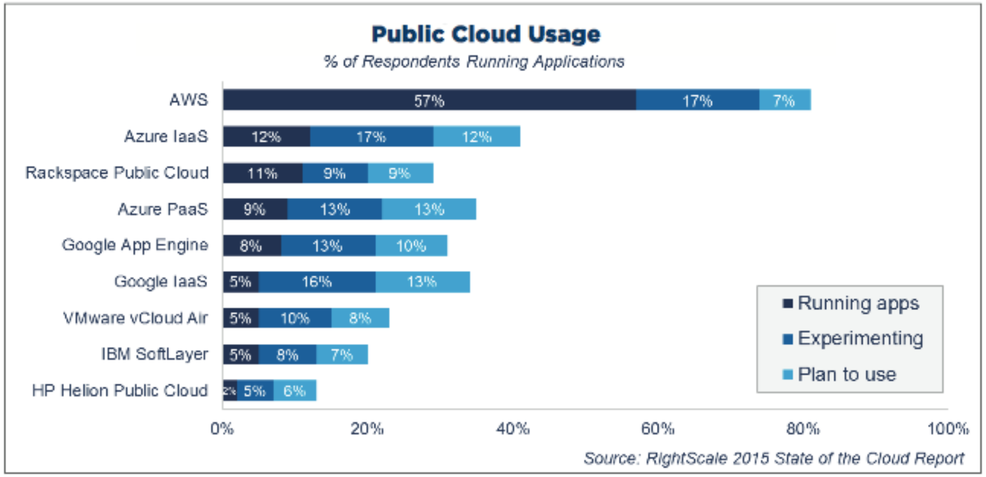

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof.

- Integration challenges that need to be managed to run at scale
  - Performance
  - Networking
  - **Provisioning**
  - Storage
  - On-demand Services
  - Monitoring and Accounting
  - Cost containment

See Steve Timm's Poster - "Virtual Machine Provisioning, Code Management and Data Movement Design for the Fermilab HEPCloud Facility" – Thu Oct 13, 11 am

Integration Challenges: Provisioning

- AWS has a **fixed price** per hour (rates vary by machine type)
- **Excess capacity** is released to the free ("spot") market **at a fraction** of the on-demand price
  - End user chooses a bid price and pays the market price. If price too high → eviction
- The **Decision Engine oversees the costs** and optimizing VM placement using the status of the facility, the historical prices, and the job characteristics.

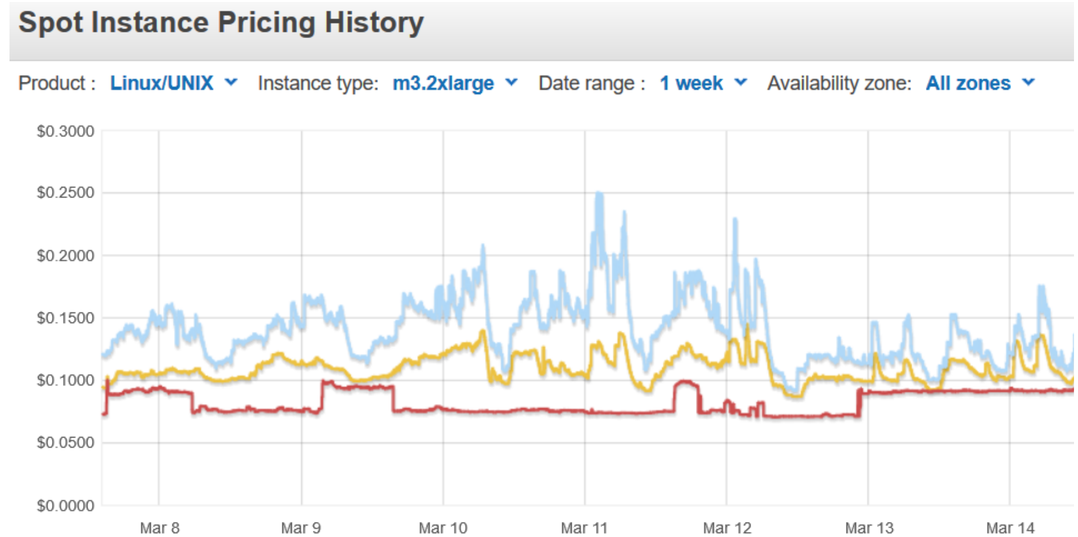

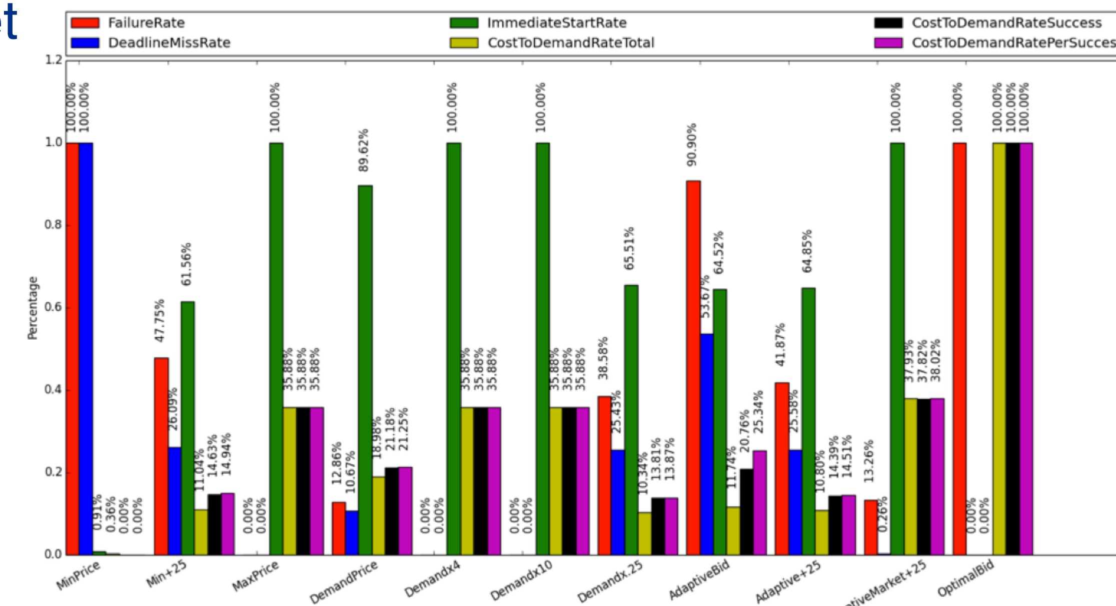

Initial HEPCloud Use Cases

NoVA Processing

Processing the 2014/2015 dataset

3 use cases: Particle ID, Montecarlo , Data Reconstruction

Received AWS research grant

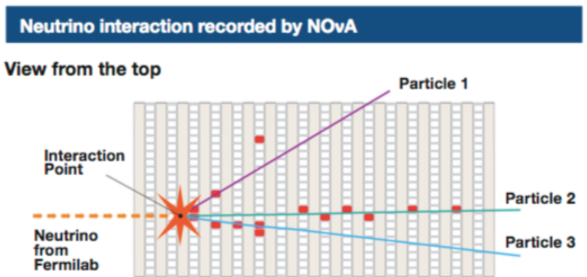

Dark Energy Survey Gravitational Waves

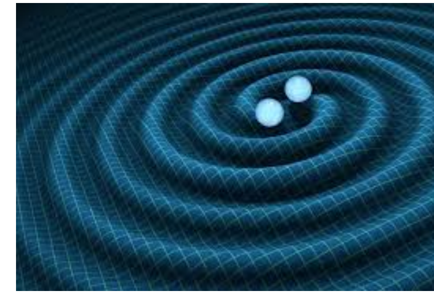

Search for optical counterpart of events detected by LIGO/VIRGO gravitational wave detectors (FNAL LDRD)

Modest CPU needs, but want 5-10 hour turnaround

Burst activity driven entirely by physical phenomena (gravitational wave events are transient)

Rapid provisioning to peak

CMS Monte Carlo Simulation

Generation (and detector simulation, digitization, reconstruction) **of simulated events** in time for **Moriond** conference.

58,000 compute cores, steady-state

Demonstrates scalability

Received AWS reasearch grant

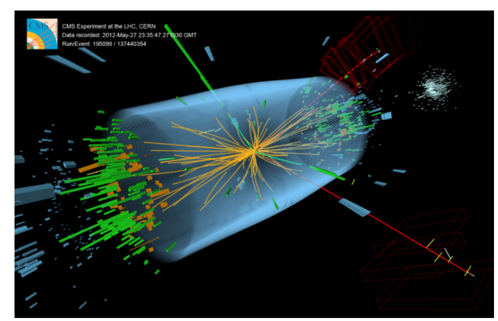

Fermilab

CMS Reaching ~60k slots on AWS with HEPCloud

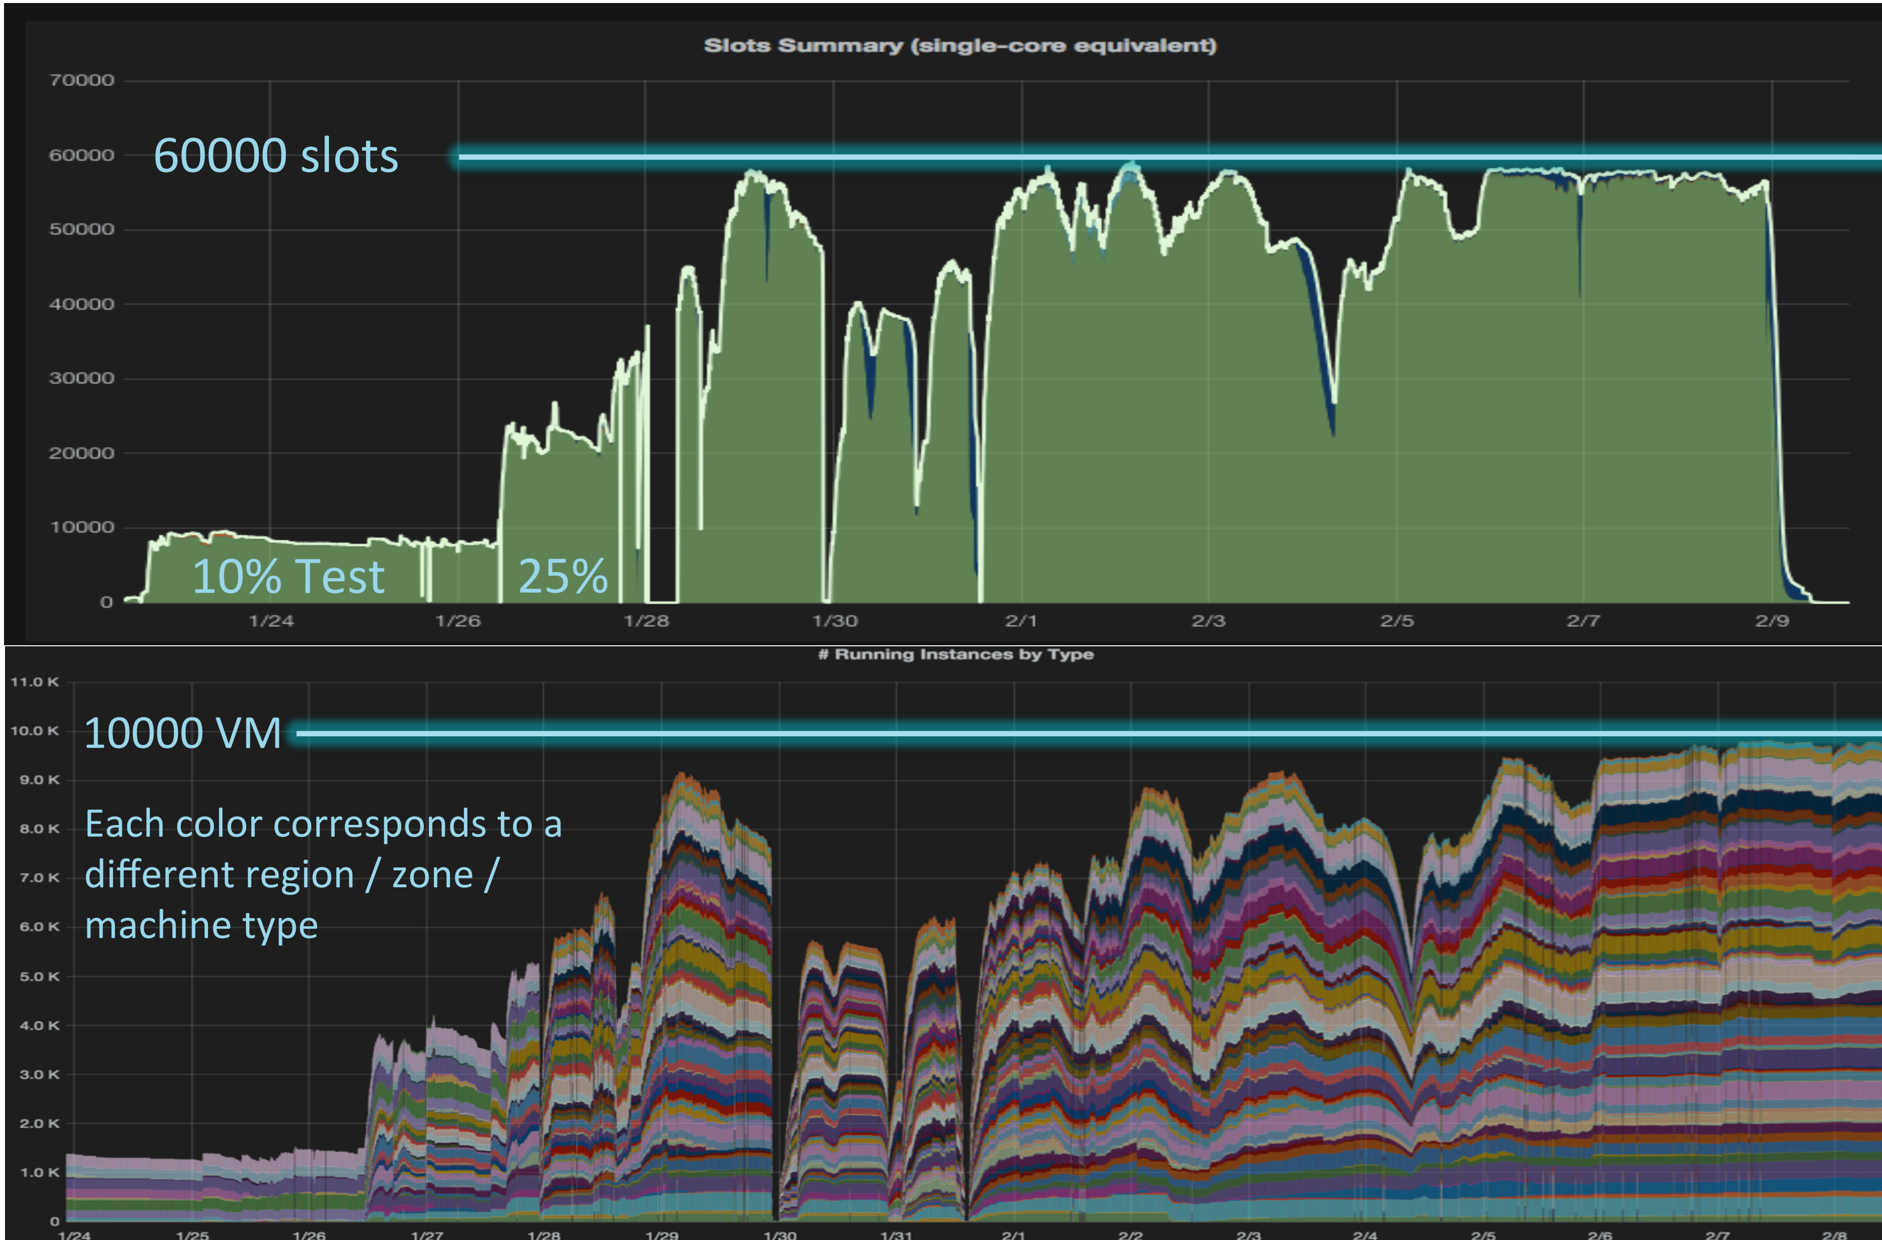

See Maria Girone's Talk - "Experience in using commercial clouds in CMS" – Thu Oct 13, 11 am

NOvA Campaigns on HEP Cloud

- Run 3 computational campaigns with different physics goals with an incremental level of complexity / improved services
  - Particle Identification → MC generation → MC Reconstruction
- **MC Reconstruction** was used to **evaluate performance and cost** of HEP Cloud
  - Input data 57 TB for 57,000 files and 114 million events
  - Processing 4.5 h per file for a total of 260,000 hours.
  - Output produced 124 TB
  - This was more data intensive than the CMS use case.
- Operations
  - Integrated the experiment Data Management system (SAM)
    - Prestage input dataset to AWS S3
    - Output is stored to S3 then asynchronously transferred to FNAL
  - Submit jobs via **usual submission interface**, then routed to HEPCloud
  - Preempted jobs automatically resubmitted
  - Resubmit jobs to reprocess files that were not fully processed because of preemption

NOvA bursts to 4x their Fermilab allocation on AWS

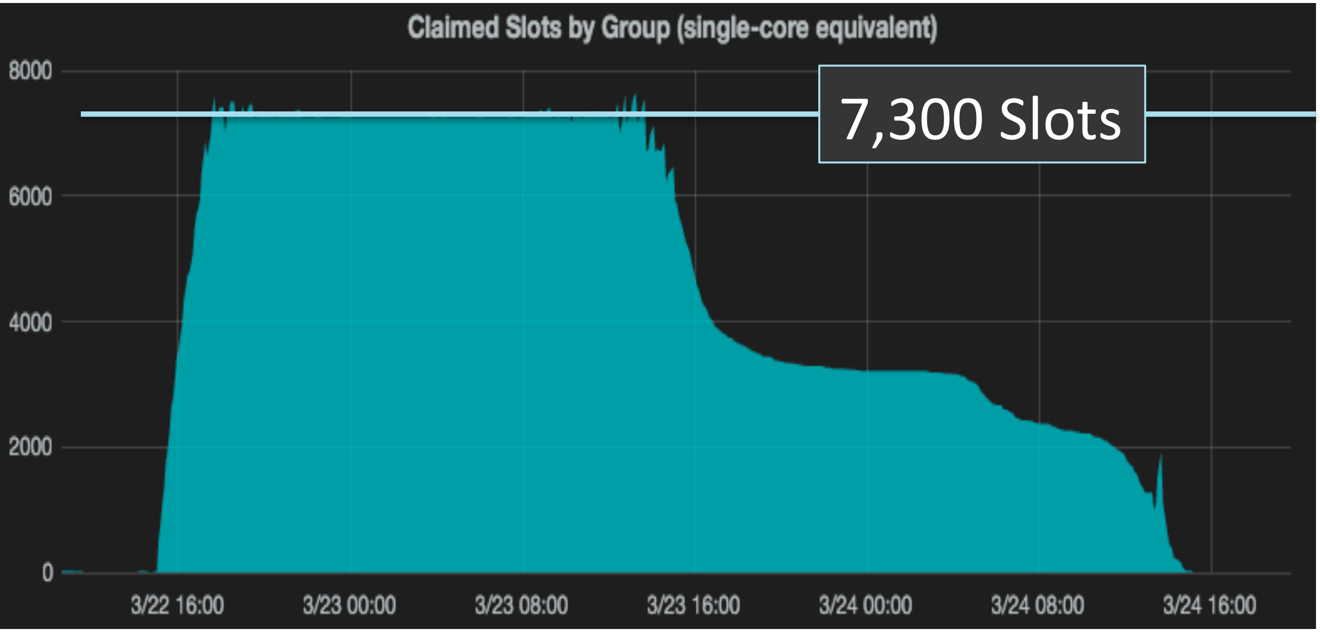

- Run ~7,300 jobs on 900 VMs for 21h out of 48 h for 203k h in one batch submission
- 10,000 jobs with 5 files per job

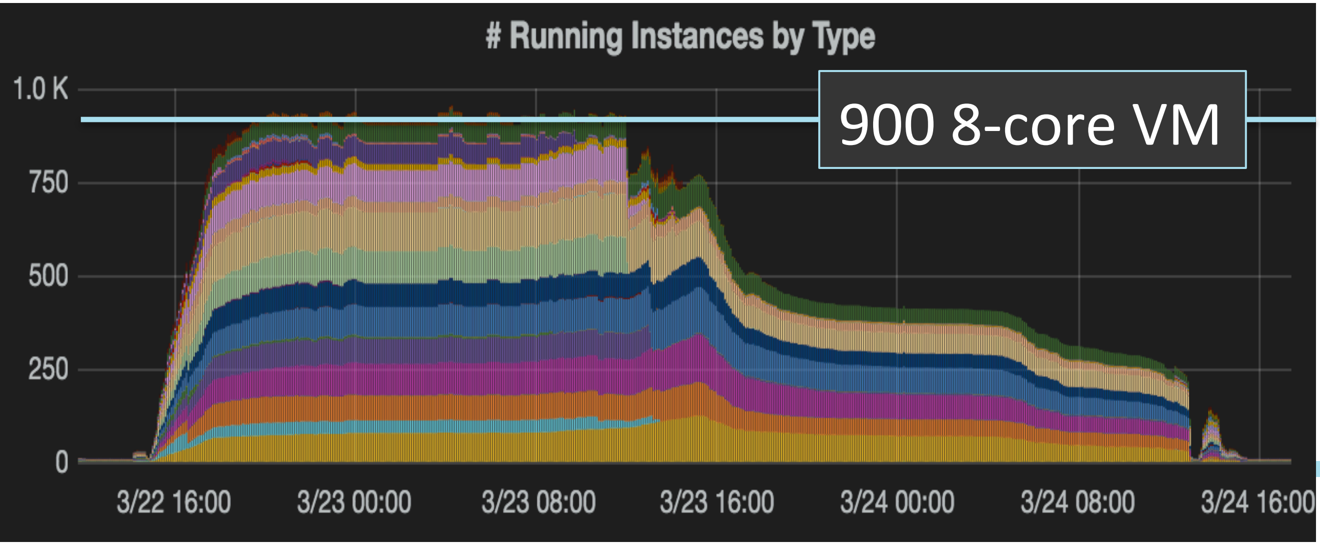

- ~6 types of 8-cores VMs in 8 overall Availability Zones in 3 AWS Regions

NOvA Costs

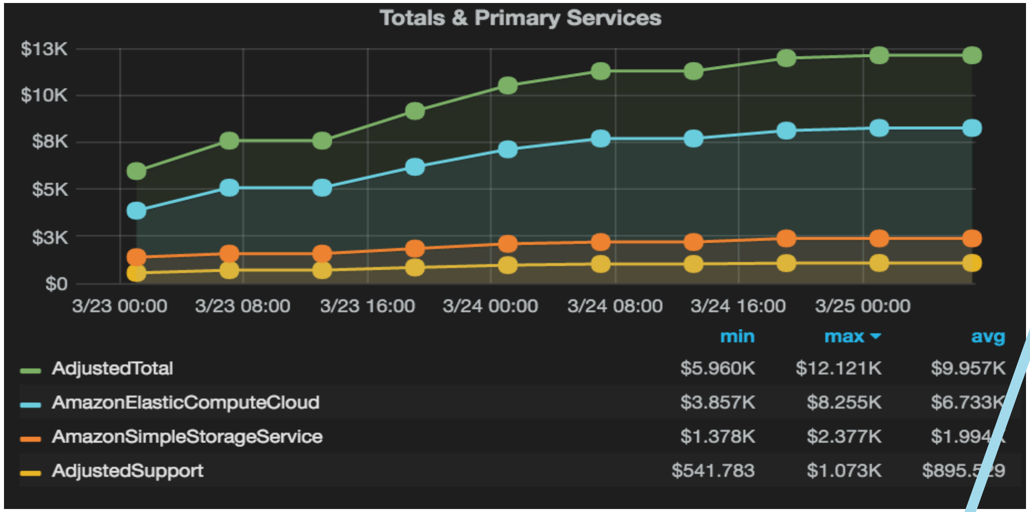

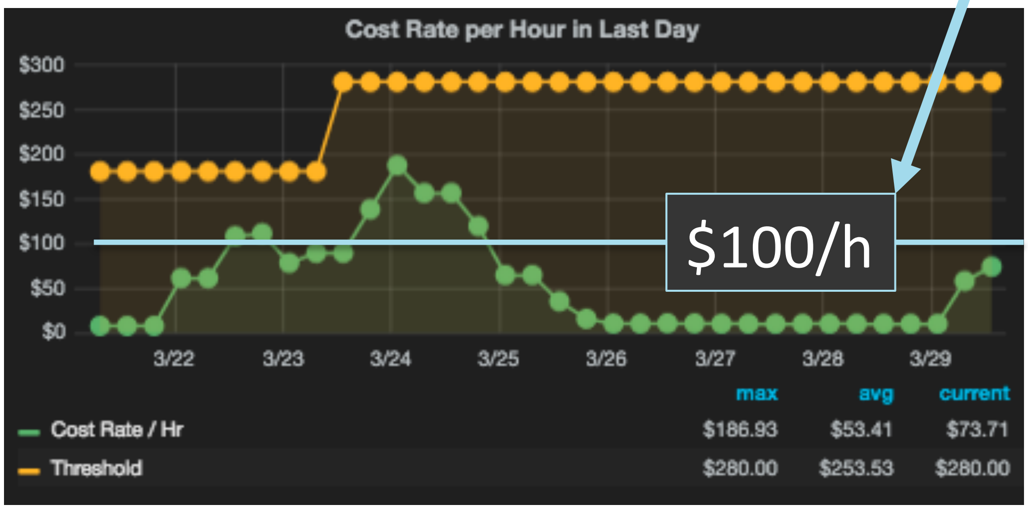

- Total Cost = $6160 (**203k h**)
  - EC2 = $4,400
  - S3 = $1,000
  - AWS Support = $530
- Approx. Rate = $100 / h @ 7k Slots
- **NO Data Egress Costs…**
  - Data egress waiver in place
  - Estimated egress (124TB) ~$7,850
- **Overall cost per consumed file (2k events) = $0.20**
- **Overall cost per core h = $0.03**
- Overhead costs:
  - …of workflow failures = $120 (jobs fail quickly)
  - …of preemption = $2,100

On-premises vs. cloud cost comparison

- Average cost per core-hour
  - On-premises resource: **0.9** cents per core-hour
    - Includes power, cooling, staff, but assumes 100% utilization
  - Off-premises at AWS (CMS use case): **1.4** cents per core-hour
  - Off-premises at AWS (NOvA use case): **3.0** cents per core-hour
    - Requirements lead to “bigger” VMs

- Benchmarks
  - Specialized (“ttbar”) benchmark focused on HEP workflows
    - On-premises: **0.163** ttbar /s (higher = better)
    - Off-premises: **0.158** ttbar /s

- Raw compute performance roughly equivalent
- Cloud costs larger – but approaching equivalence

Acknowledgements

- The support from the Computing Sector
- The Fermilab HEPCloud Facility team
- AWS and their engagement team, in particular Jamie Baker
- The HTCondor team
- The collaboration and contributions from KISTI, in particular Dr. Seo-Young Noh
- The Illinois Institute of Technology (IIT) students and professors Ioan Raicu and Shangping Ren
- The Italian National Institute of Nuclear Physics (INFN) summer student program

Backup

HEPCloud Compute and HPC

- A very appealing possibility, as we are approaching the exascale era, is to consider HPC facilities as a potential compute resource for HEPCloud
  - and, in the other direction, consider HEPCloud facility services (e.g. storage) as a potential resource for HPC facilities
- Investigate use cases with workflows that will allow such utilization within the constraints of allocation, security and access policy of HPC facilities.
- Initiate work with HPC facilities to fully understand constraints and requirements that will enable us to develop the HEPCloud process, policies and tools necessary for access of HPC resources

NOvA: Neutrino Experiment

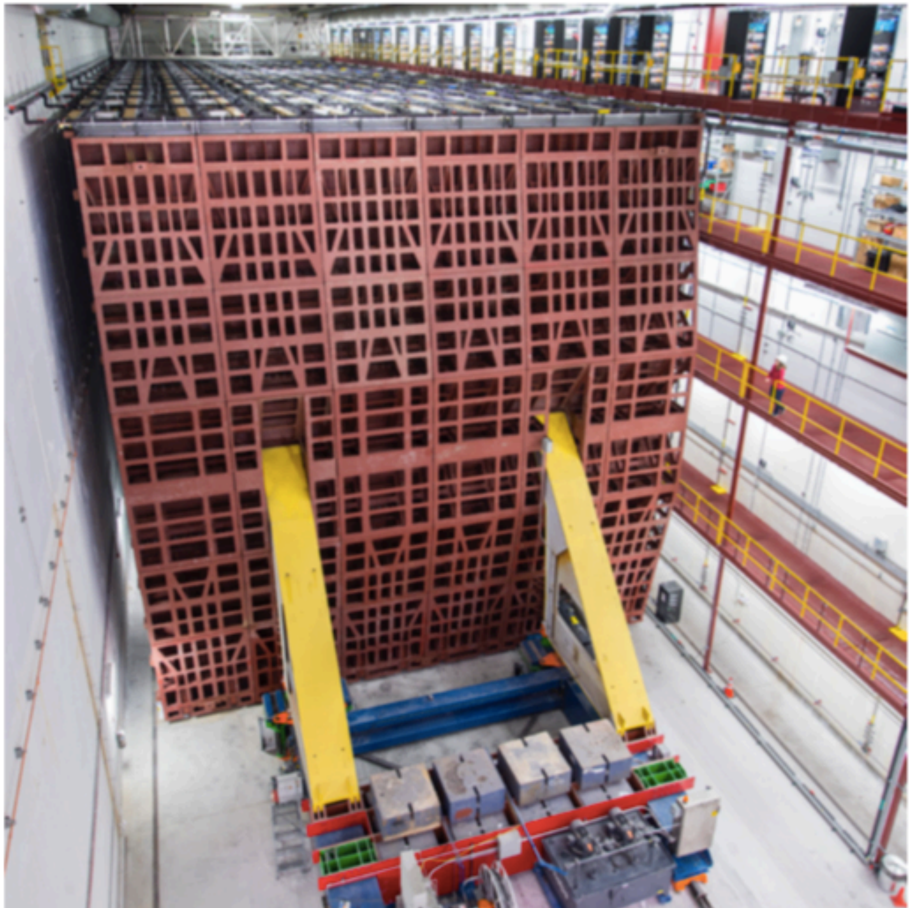

The NOvA detector in Minnesota occupies an area about the size of two basketball courts. It is 200 feet long and made of modules 50 feet high and 50 feet wide. The detector records particle tracks from neutrinos sent by a powerful accelerator at Fermilab. The construction of the NOvA detectors was completed in the fall of 2014, on time and under budget. The experiment is scheduled to collect information for six years.

Neutrino interaction recorded by NOvA

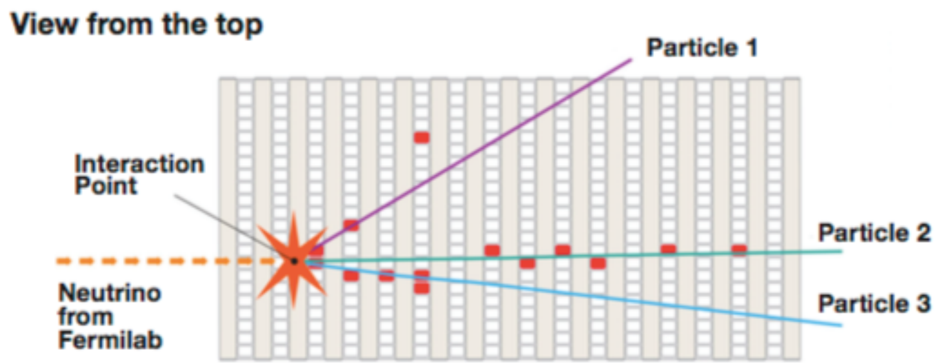

Neutrinos rarely interact with matter. When a neutrino smashes into an atom in the NOvA detector in Minnesota, it creates distinctive particle tracks. Scientists explore these particle interactions to better understand the transition of muon neutrinos into electron neutrinos. The experiment also helps answer important scientific questions about neutrino masses, neutrino oscillations, and the role neutrinos played in the early universe.

NOvA Use Case

NoVA Processing

Processing the 2014/2015 dataset
16 4-day "campaigns" over one year
Demonstrates stability, availability, cost-effectiveness
Received AWS academic grant

Dark Energy Survey - Gravitational Waves

Search for optical counterpart of events detected by LIGO/VIRGO gravitational wave detectors (FNAL LDRD)
Modest CPU needs, but want 5-10 hour turnaround
Burst activity driven entirely by physical phenomena (gravitational wave events are transient)
Rapid provisioning to peak

CMS Monte Carlo Simulation

Generation (and detector simulation, digitization, reconstruction) of simulated events in time for Moriond conference
56000 compute cores, steady-state
Demonstrates scalability
Received AWS academic grant

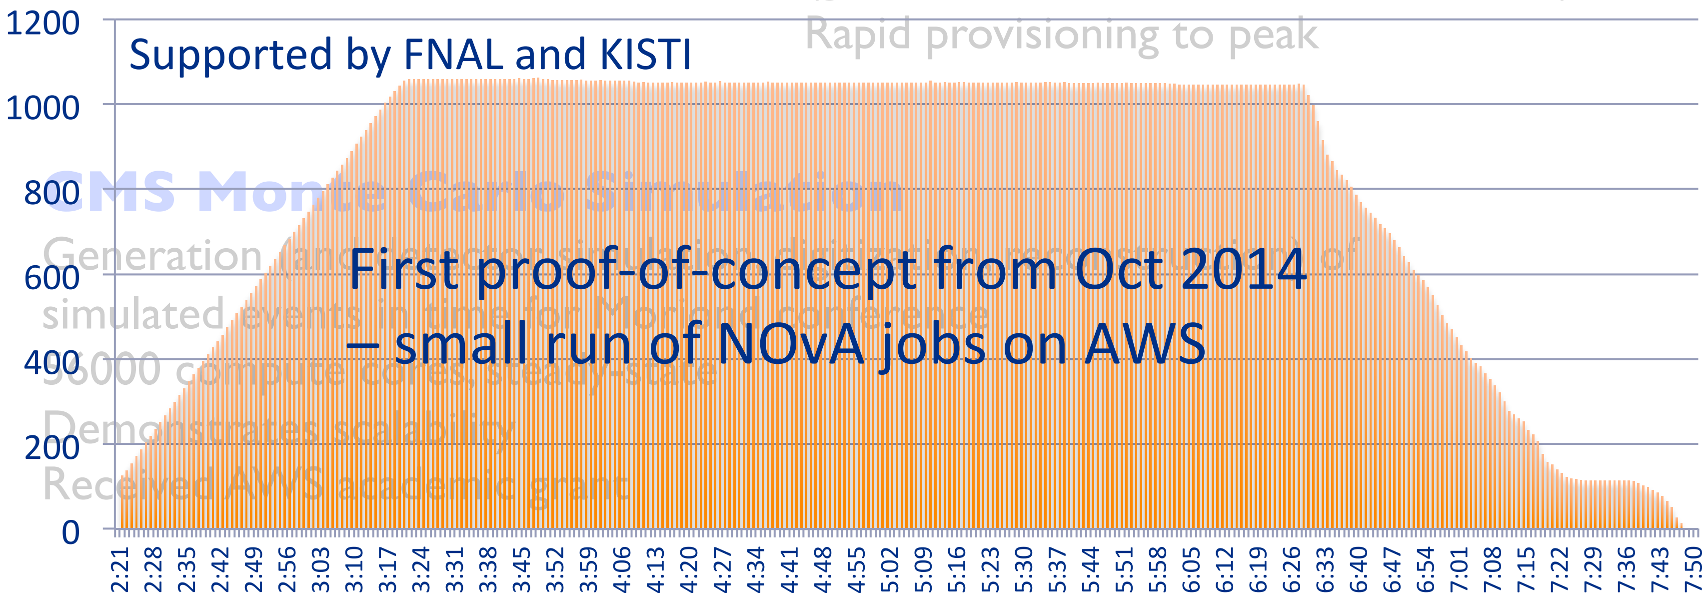

Results from the CMS Use Case

- **All CMS simulation requests fulfilled for Moriond**
  - **2.9 million jobs, 15.1 million wall hours**
    - 9.5% badput – includes preemption from spot pricing
    - 87% CPU efficiency
  - **518 million events generated**

/DYJetsToLL_M-50_TuneCUETP8M1_13TeV-amcatnloFXFX-pythia8/RunIIFall15DR76-PU25nsData2015v1_76X_mcRun2_asymptotic_v12_ext4-v1/AODSIM
/DYJetsToLL_M-10to50_TuneCUETP8M1_13TeV-amcatnloFXFX-pythia8/RunIIFall15DR76-PU25nsData2015v1_76X_mcRun2_asymptotic_v12_ext3-v1/AODSIM
/TTJets_13TeV-amcatnloFXFX-pythia8/RunIIFall15DR76-PU25nsData2015v1_76X_mcRun2_asymptotic_v12_ext1-v1/AODSIM
/WJetsToLNu_TuneCUETP8M1_13TeV-amcatnloFXFX-pythia8/RunIIFall15DR76-PU25nsData2015v1_76X_mcRun2_asymptotic_v12_ext4-v1/AODSIM

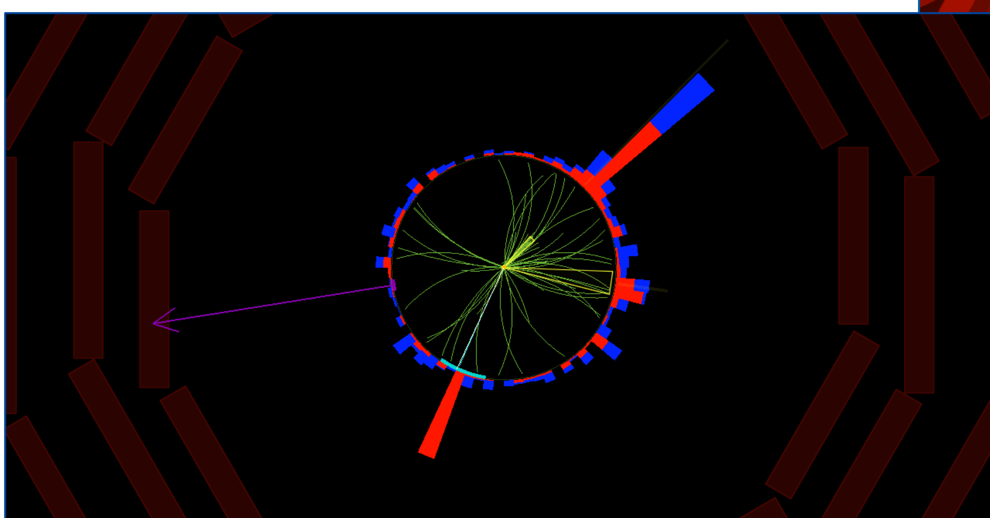

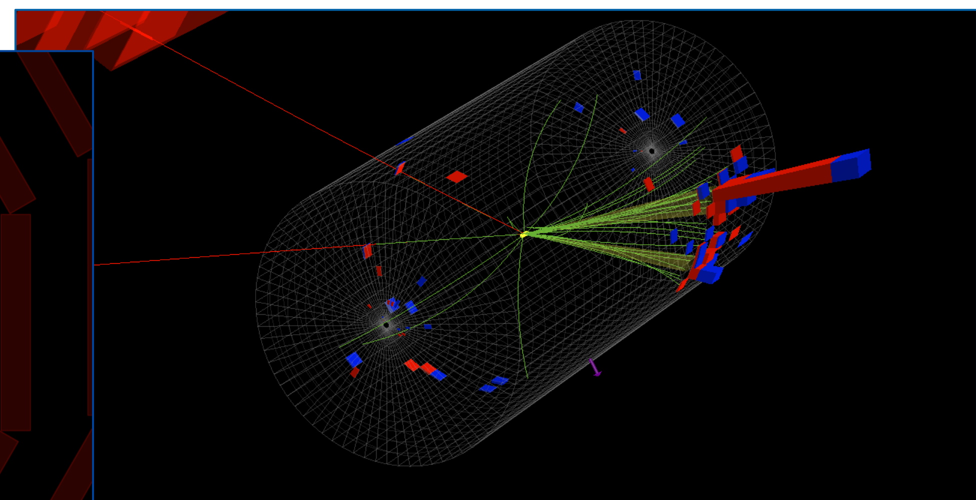

HEPCloud AWS slots by Region/Zone

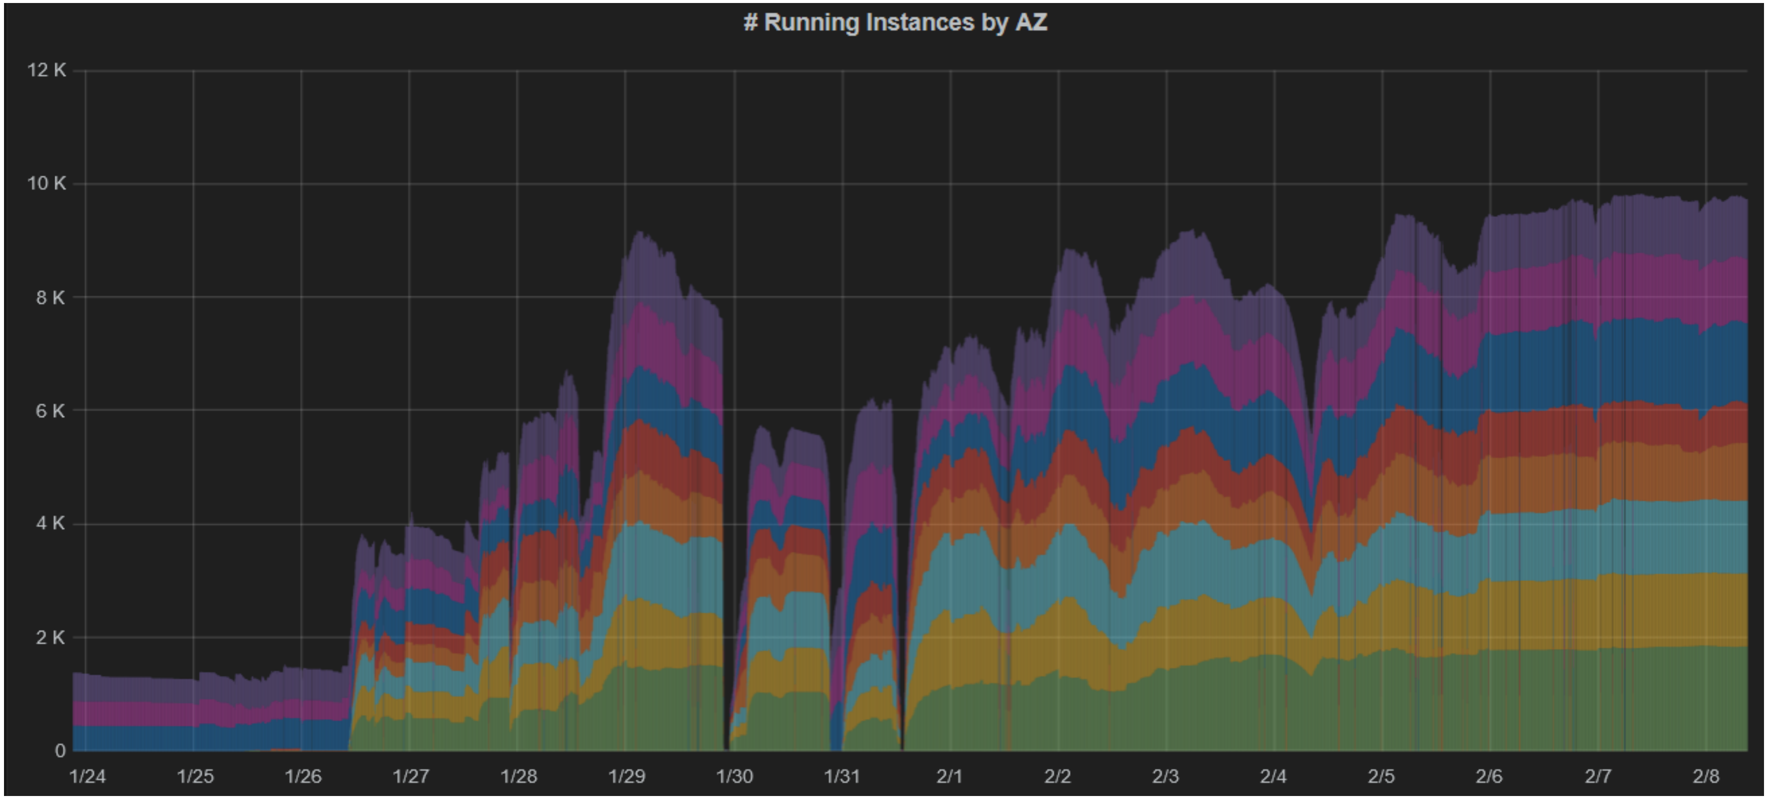

Each color corresponds to a different region+zone

HEPCloud AWS slots by Region/Zone/Type

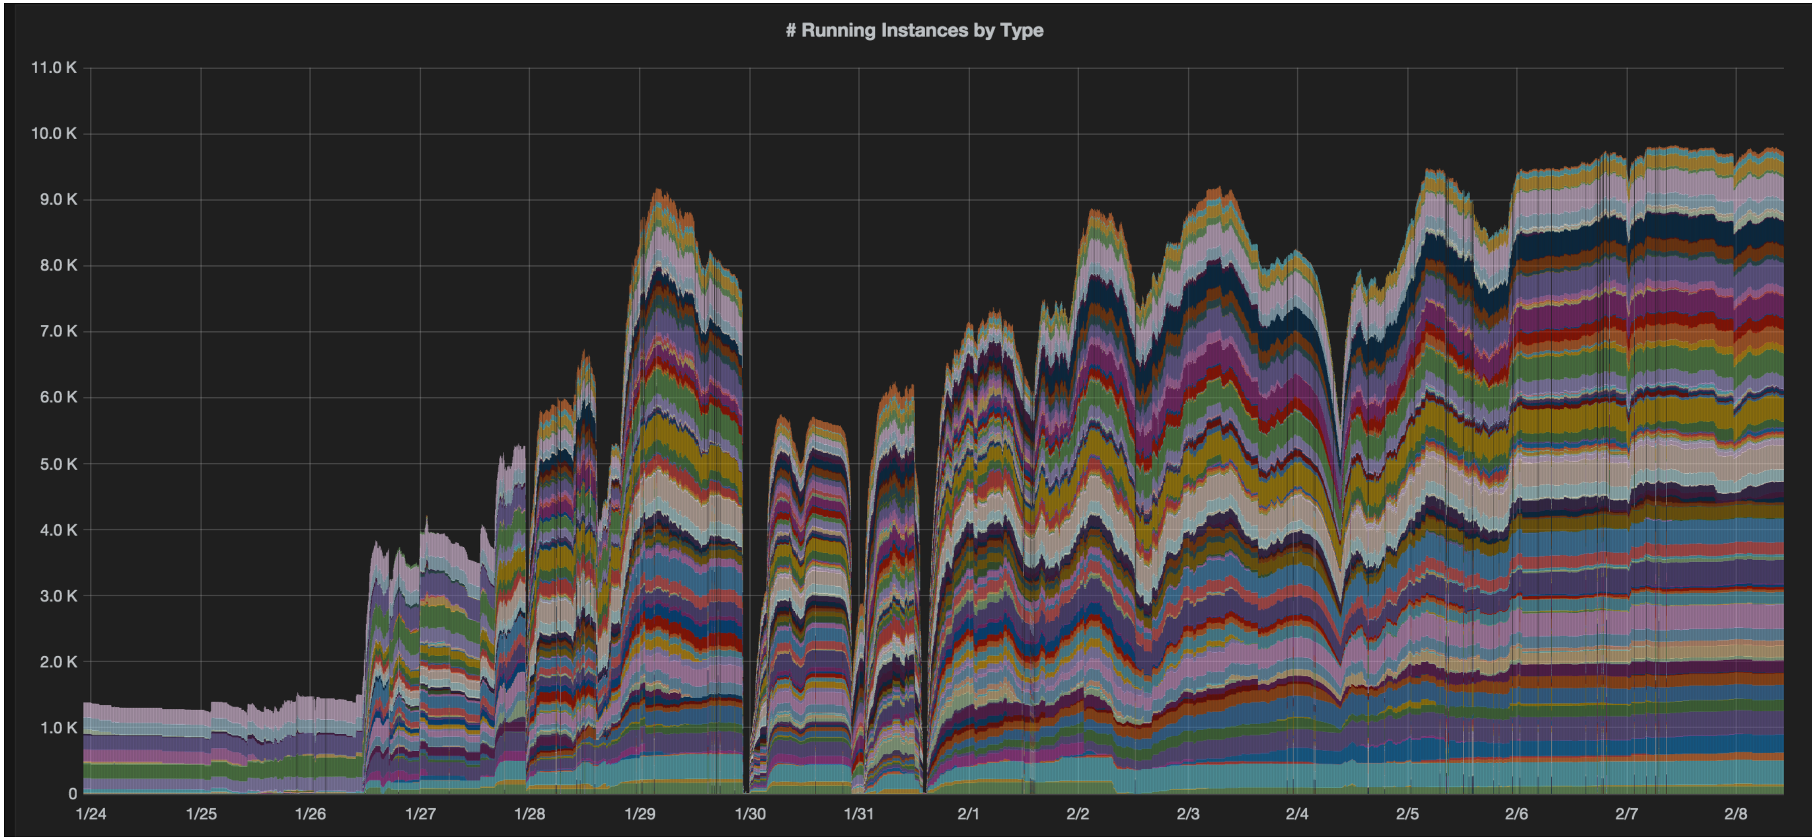

Each color corresponds to a different region+zone+machine type

HEPCloud AWS: 25% of CMS global capacity

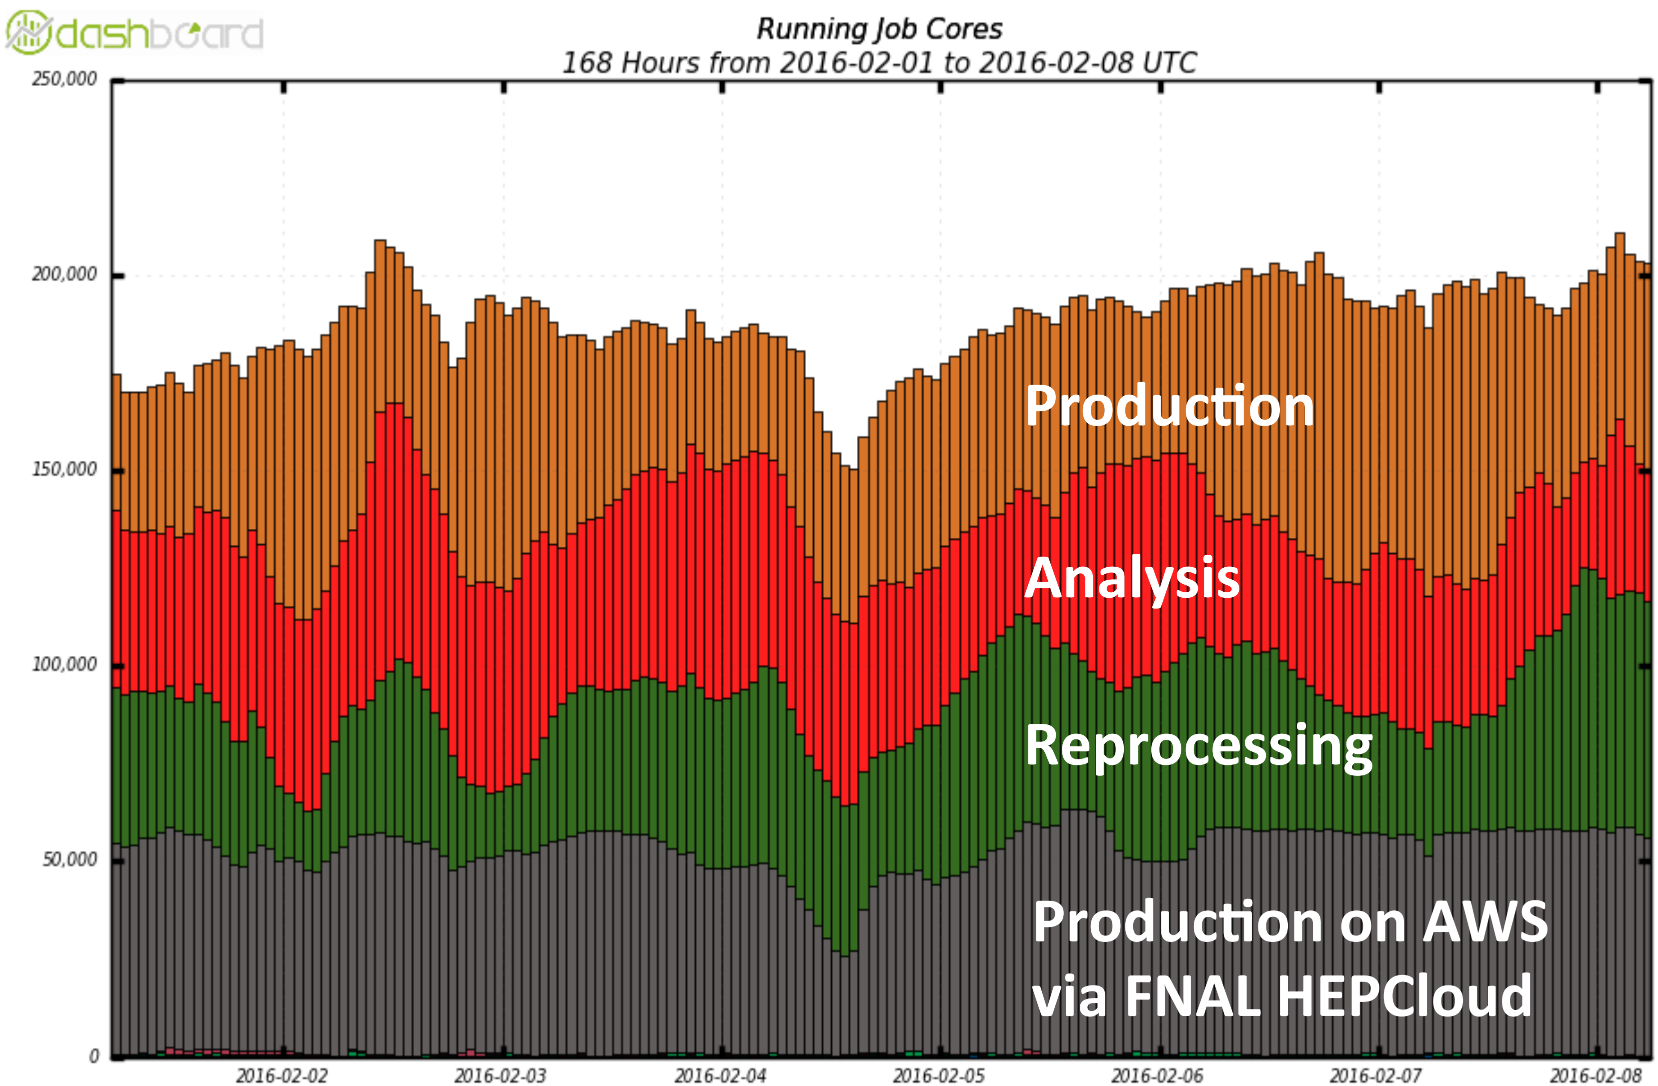

Fermilab HEPCloud compared to global CMS Tier-1

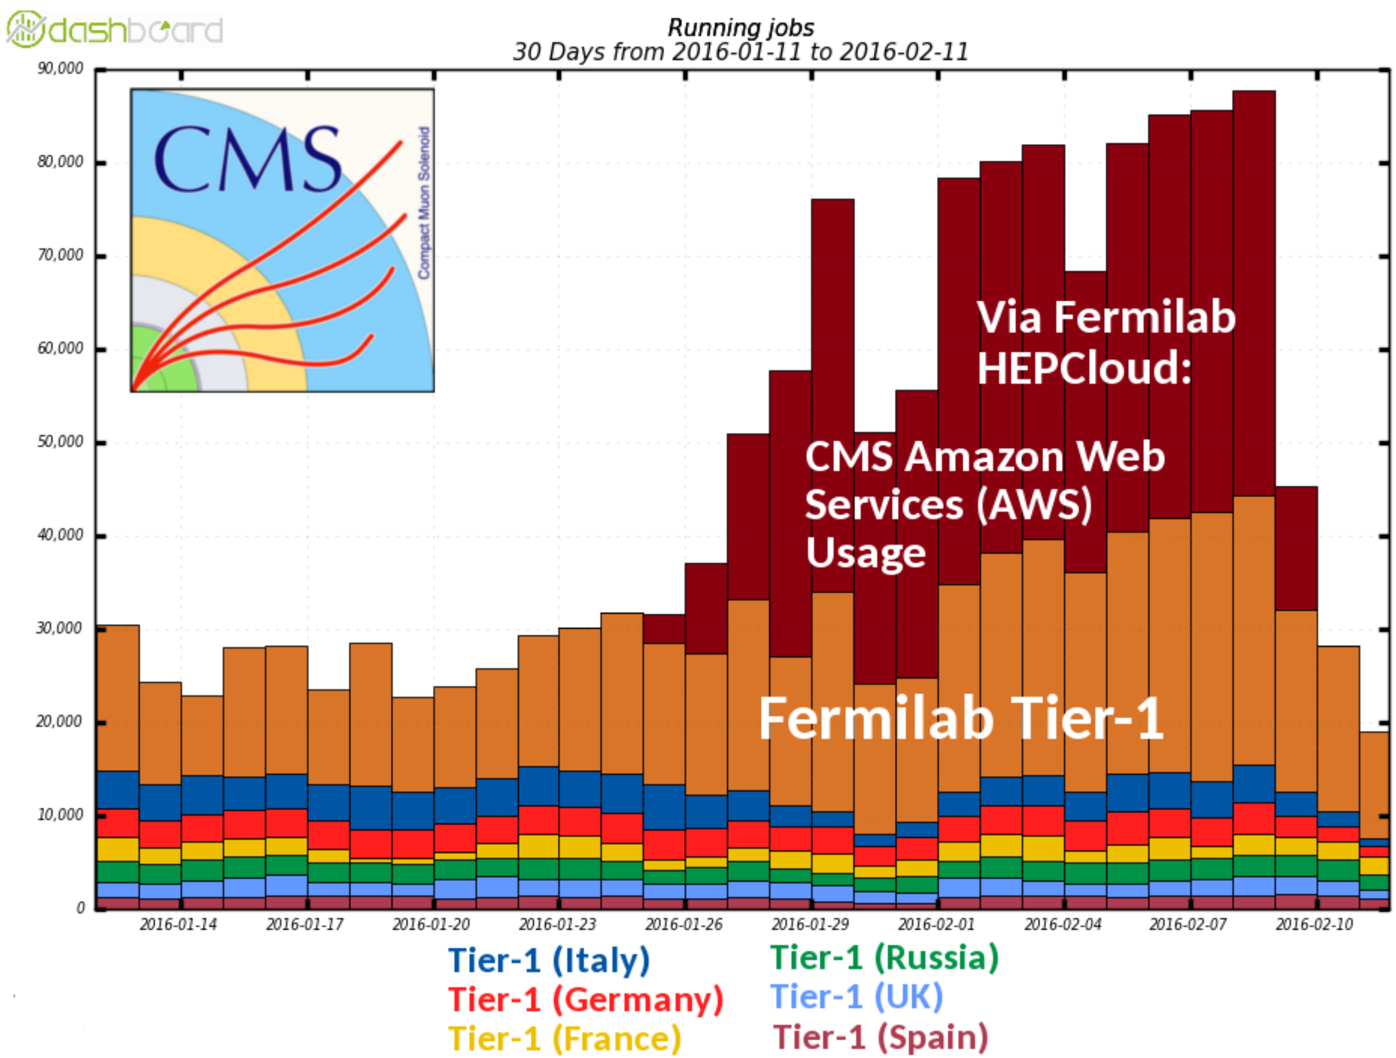

HEPCloud: Orchestration

- Monitoring and Accounting
  - Synergies with FIFE monitoring projects
    - But also monitoring real-time expense
  - Feedback loop into Decision Engine

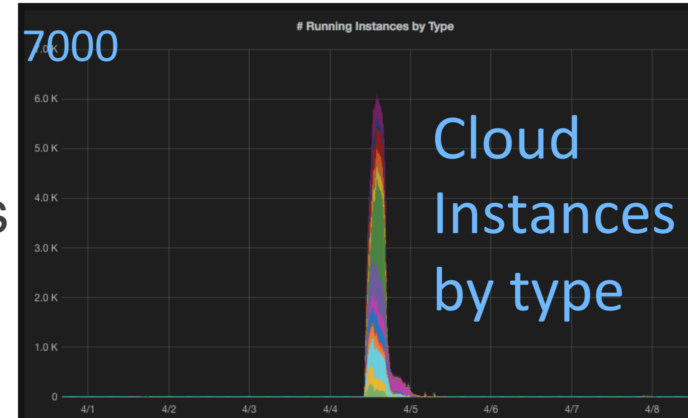

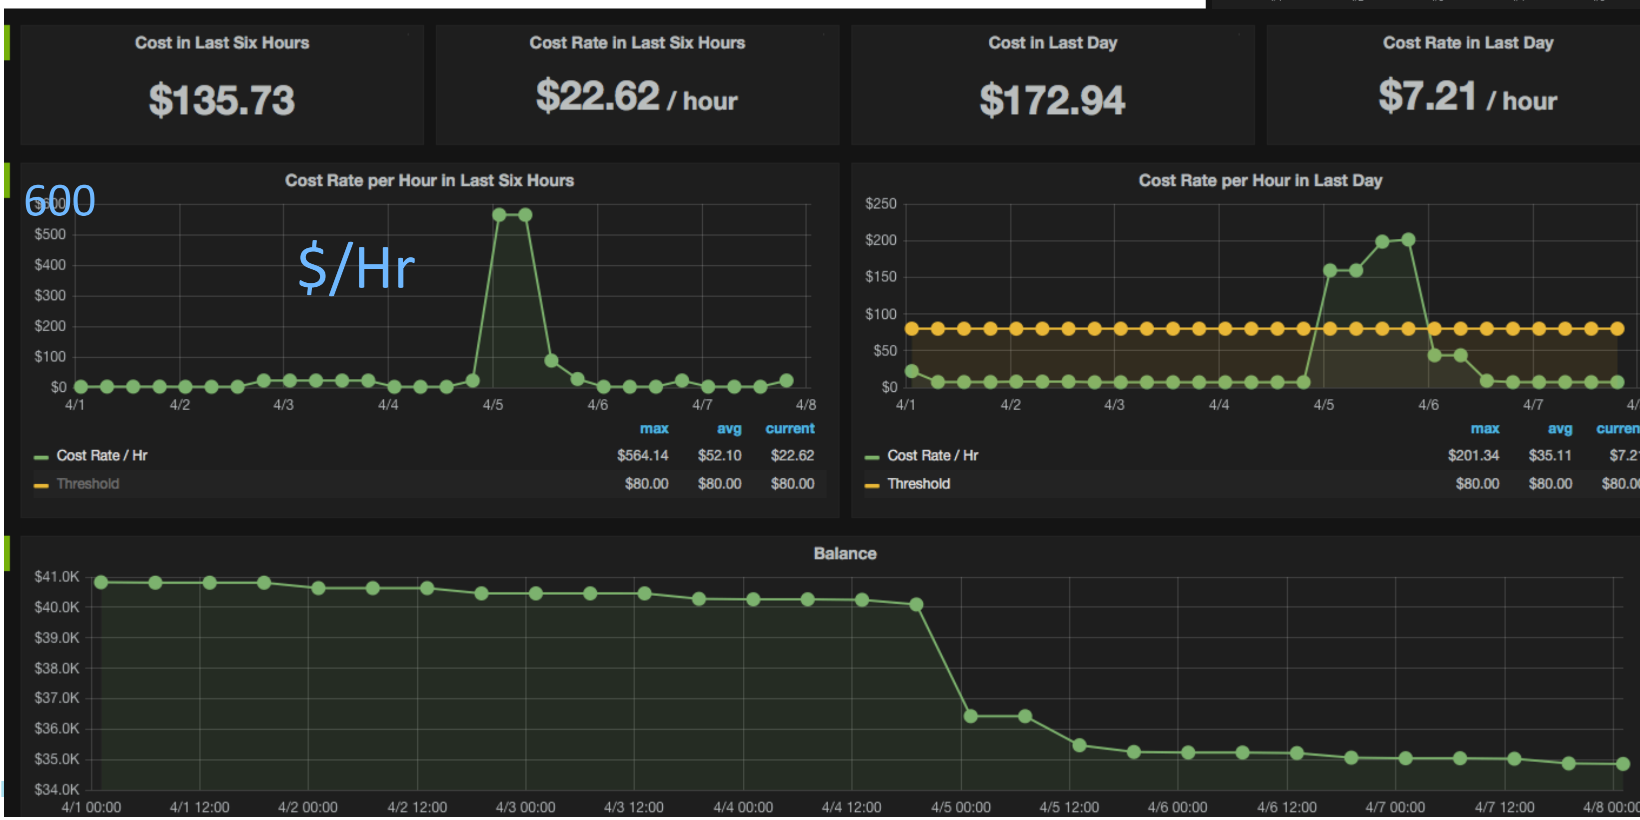

For More Information:

- NOvA: http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=5774
- CMS: http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=5750
- MicroBooNE@NERSC: http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=5693

HEPCloud Collaborations

- Engage in collaboration to leverage tools and experience whenever possible

- Grid technologies – Worldwide LHC Computing Grid
  - Preparing communities for distributed computing
- BNL and ATLAS – engaged in next HEPCloud phase
- HTCondor – common provisioning interface (underneath Panda, glideinWMS)
- CMS – collaborative knowledge and tools, cloud-capable workflows
- CERN faces similar challenges and we are having productive conversations with different facets
  - For example - CERN openlab CTO is engaged in HEPCloud

Decision Engine – design & architecture

- Decision Engine chooses what to provision next
  - v1.5 implementation: Strict matching based on processing type
  - v2.0 implementation: Zeroth-order prioritization based on cost

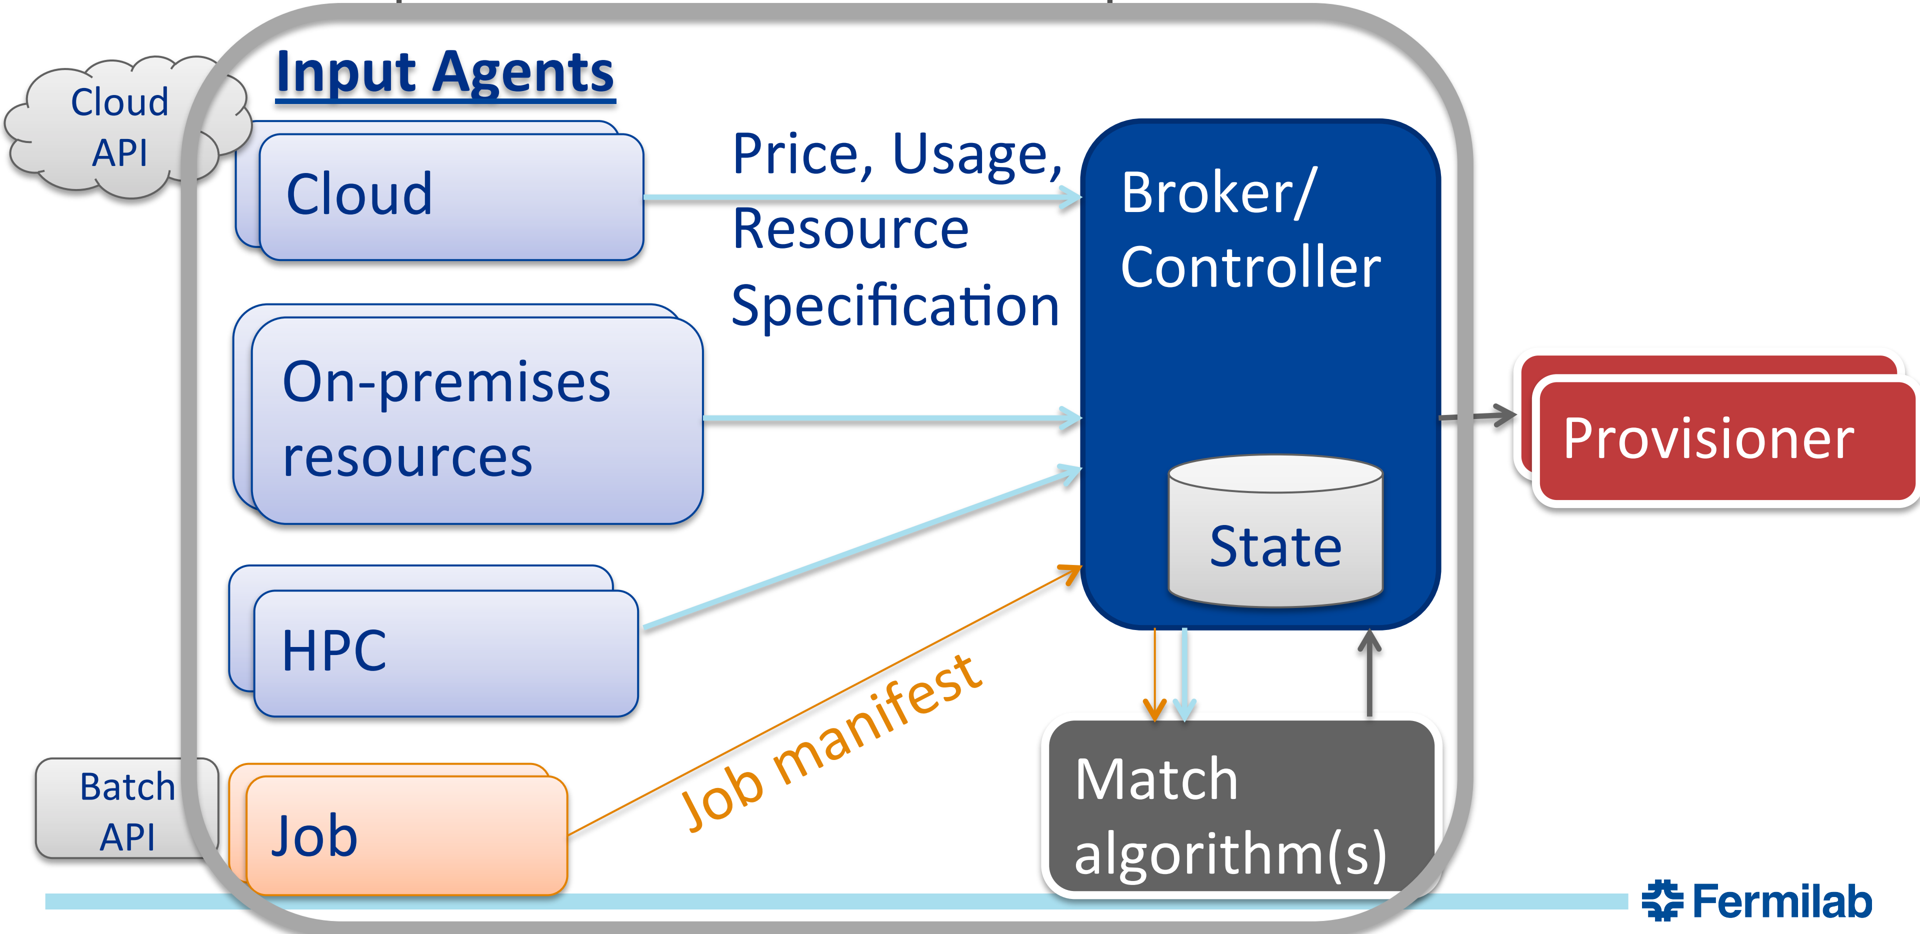

Decision Engine v2.0 Implementation

- Match algorithm calculates a single **figure of merit** per resource
- Goal
  - Spread jobs equally among resource types that have best price/performance ratio
  - Divest resources to avoid increasing the price and increasing pre-emption pressure
- Inputs
  - Performance for each resource type
  - Current "spot" price for each resource type/region/zone
  - Number of provisioned resources in each type/region/zone
- Policy
  - Bid **25%** of "on-demand" price
  - Calculate price/performance **P** for each type/region/zone
  - Calculate occupancy fraction **C=(N+1)/100*(bid_price – spot_price - .001)**
  - Calculate figure of merit **F** = **C***P** (lower = better)
- Output
  - Provision resources with **10 lowest** figures of merit

Fermilab

Pythia on Mira – Details

- We incorporated MPI into the main routines, using scatter and broadcast to send out unique parameters.   The plan is to start one process on each node, running 64 threads, each with an instance of the pythia-based analysis. We will do this in chunk of 128 nodes, where each chunk is a gather collection point for writing to disk.

- Things were running on our x86 cluster - the porting to power PC of the build tools was the challenging part.

- ~150 core test

Description of CMS workflow

- Four chained steps (output of step N is input of step N+1)
  - Step 1 requires few GB input (“Gridpack”) – same files per job
  - Step 3 requires additional input: “pile-up” data (simulating multiple events per bunch crossing), 5-10 GB

- Pile-up data is constructed on-the-fly by random seek and sequential reads into a 10 TB dataset
    - More than 150 pile-up events read per simulated event – this step is extremely I/O intensive
    - Staged pile-up datasets to AWS S3 (storage service) ahead-of-time using FTS3

Reading pile-up from AWS S3 (storage)

- AWS worker nodes granted permission to read from AWS S3 folder ("bucket") via AWS Security-Token-Service (STS)
- ROOT has a TS3WebFile class!
  - But session key support was missing (needed for STS!)

Add support for session keys in TS3WebFile

(some minor revision also by the committer)

Browse files

master v6-07-04

holzman committed with **smithdh** on Dec 1, 2015 1 parent 55ed62b commit fe169587a0dc681a33ecdd33544c32cbeb43d3b7

Showing 4 changed files with **73 additions** and **14 deletions**.

Unified Split

- This worked great!
  - Except…

Reading pile-up from AWS S3 (storage)

- Cost of data access was 30% of compute costs
  - 150 million HTTP GETs per hour is a lot!
- Wrote a curl wrapper to provide the custom AWS authentication headers
  - (Not often I can say I reduced costs by 5 orders of magnitude!)

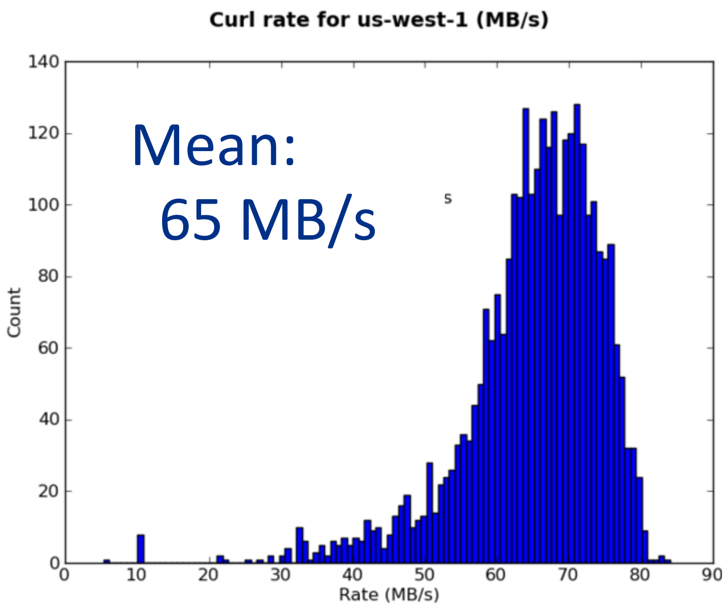

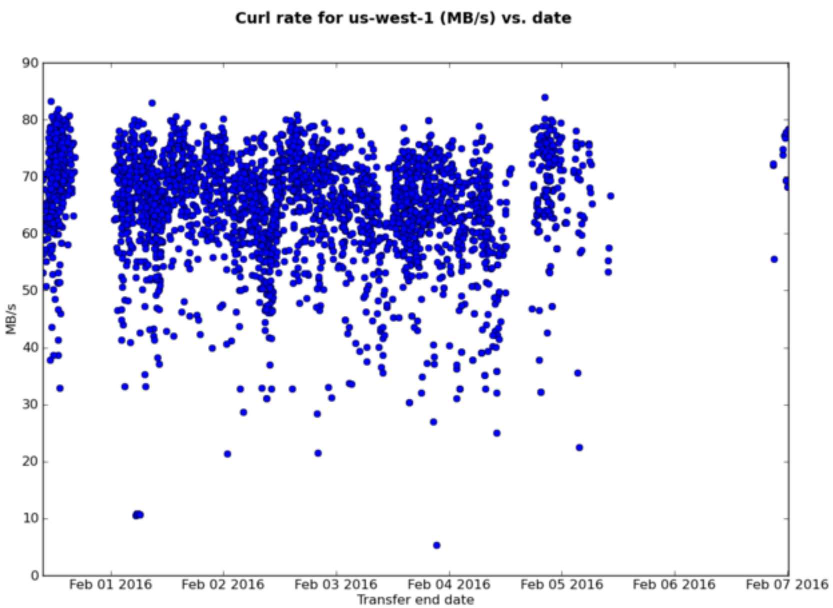

Preemption on AWS

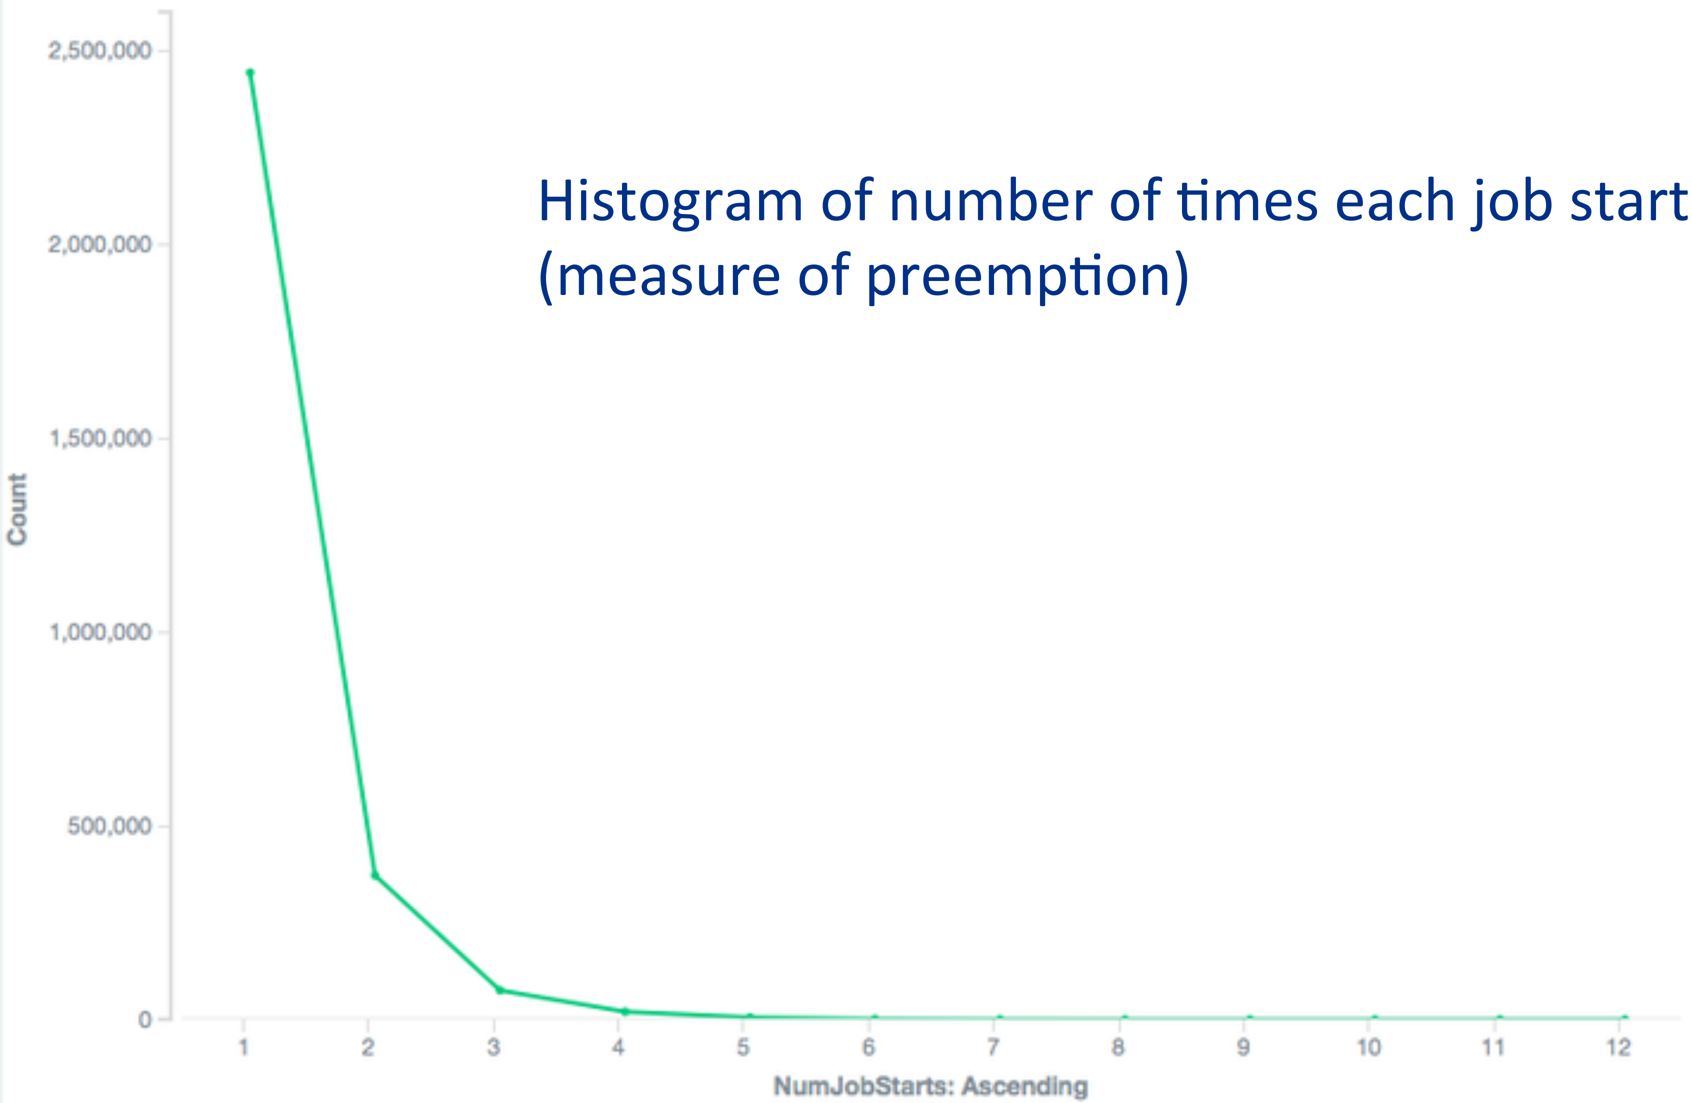

Histogram of number of times each job started (measure of preemption)

AWS: some instances are more equal than others

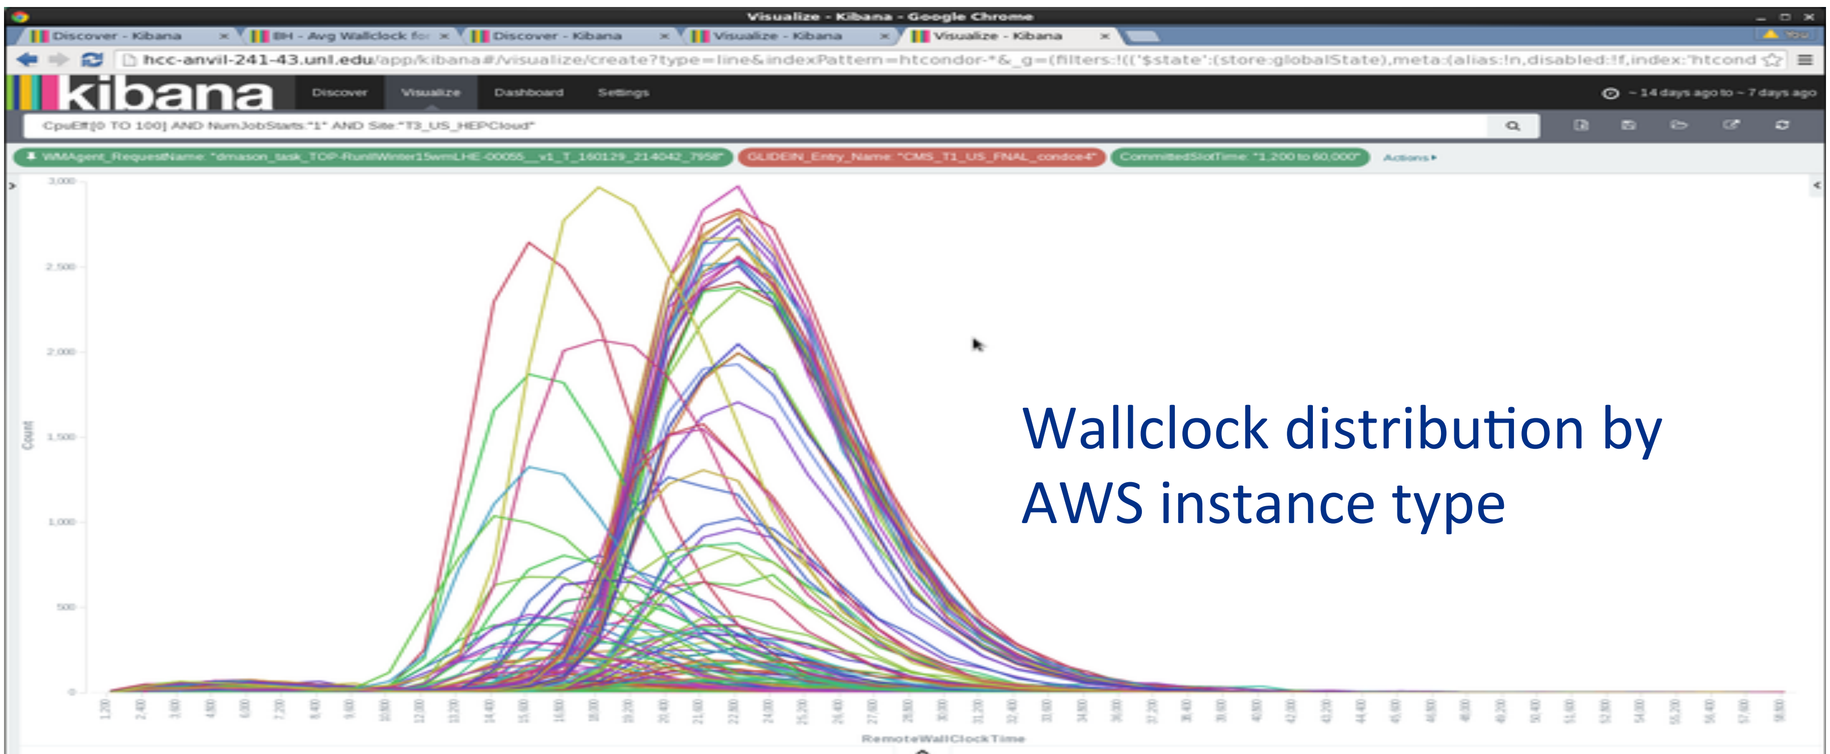

Elements of the cost per core-hour

Based on Fermilab CMS Tier-1

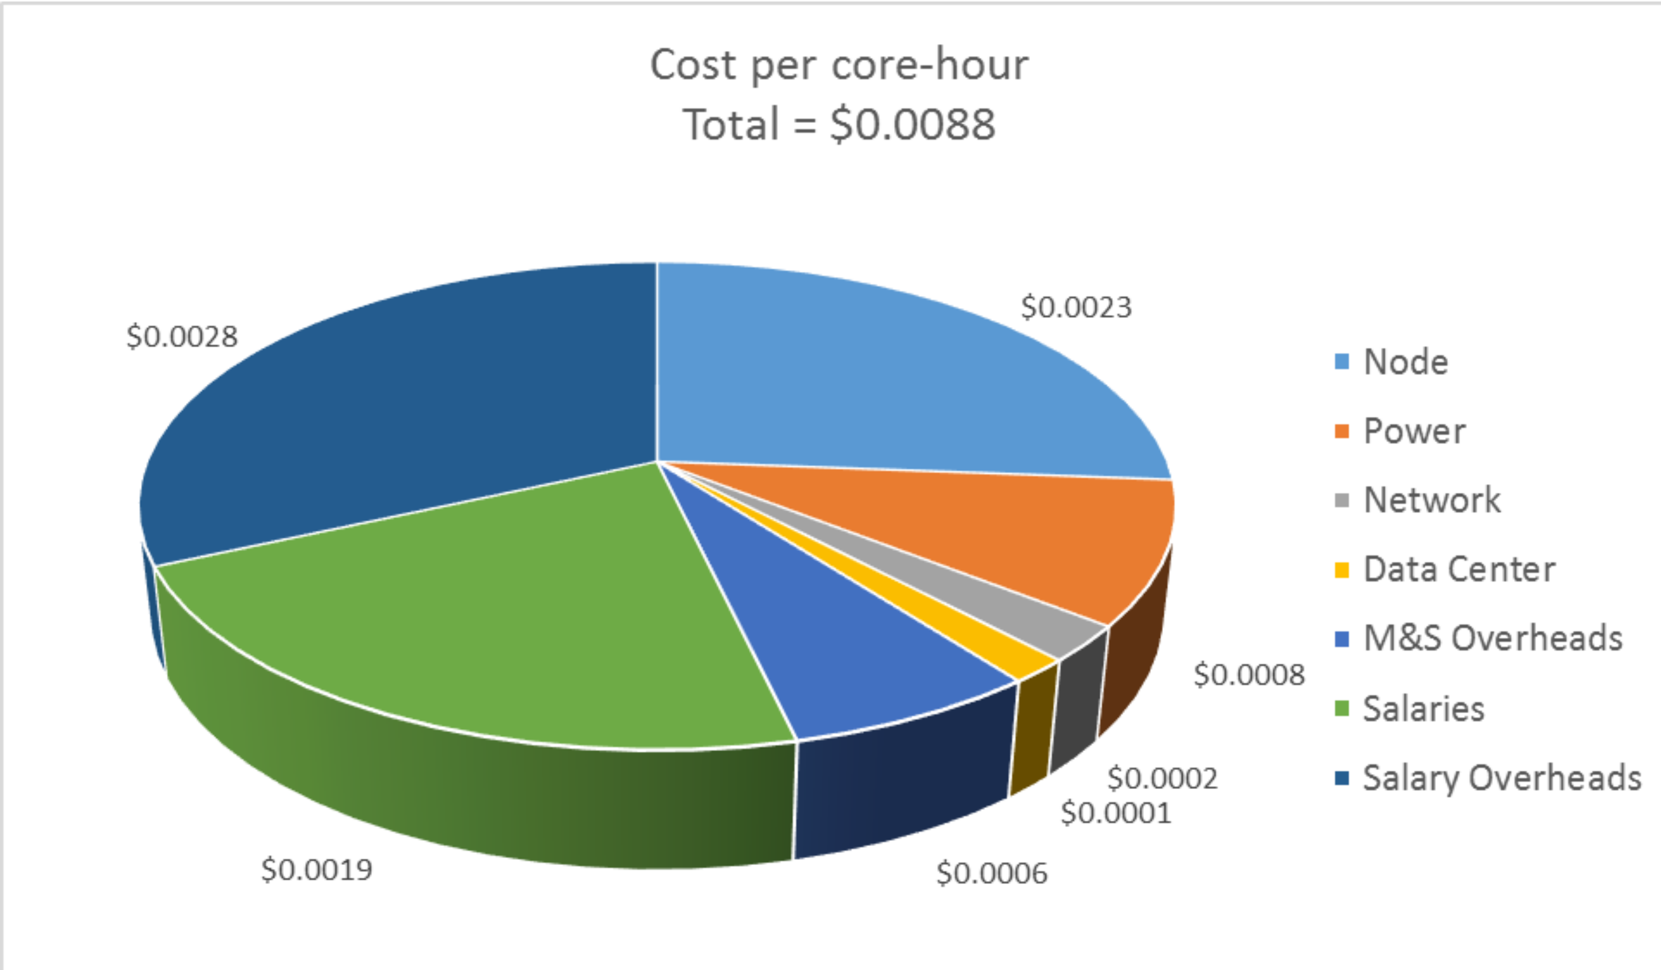

glideinWMS – Building dynamic HTCondor pools

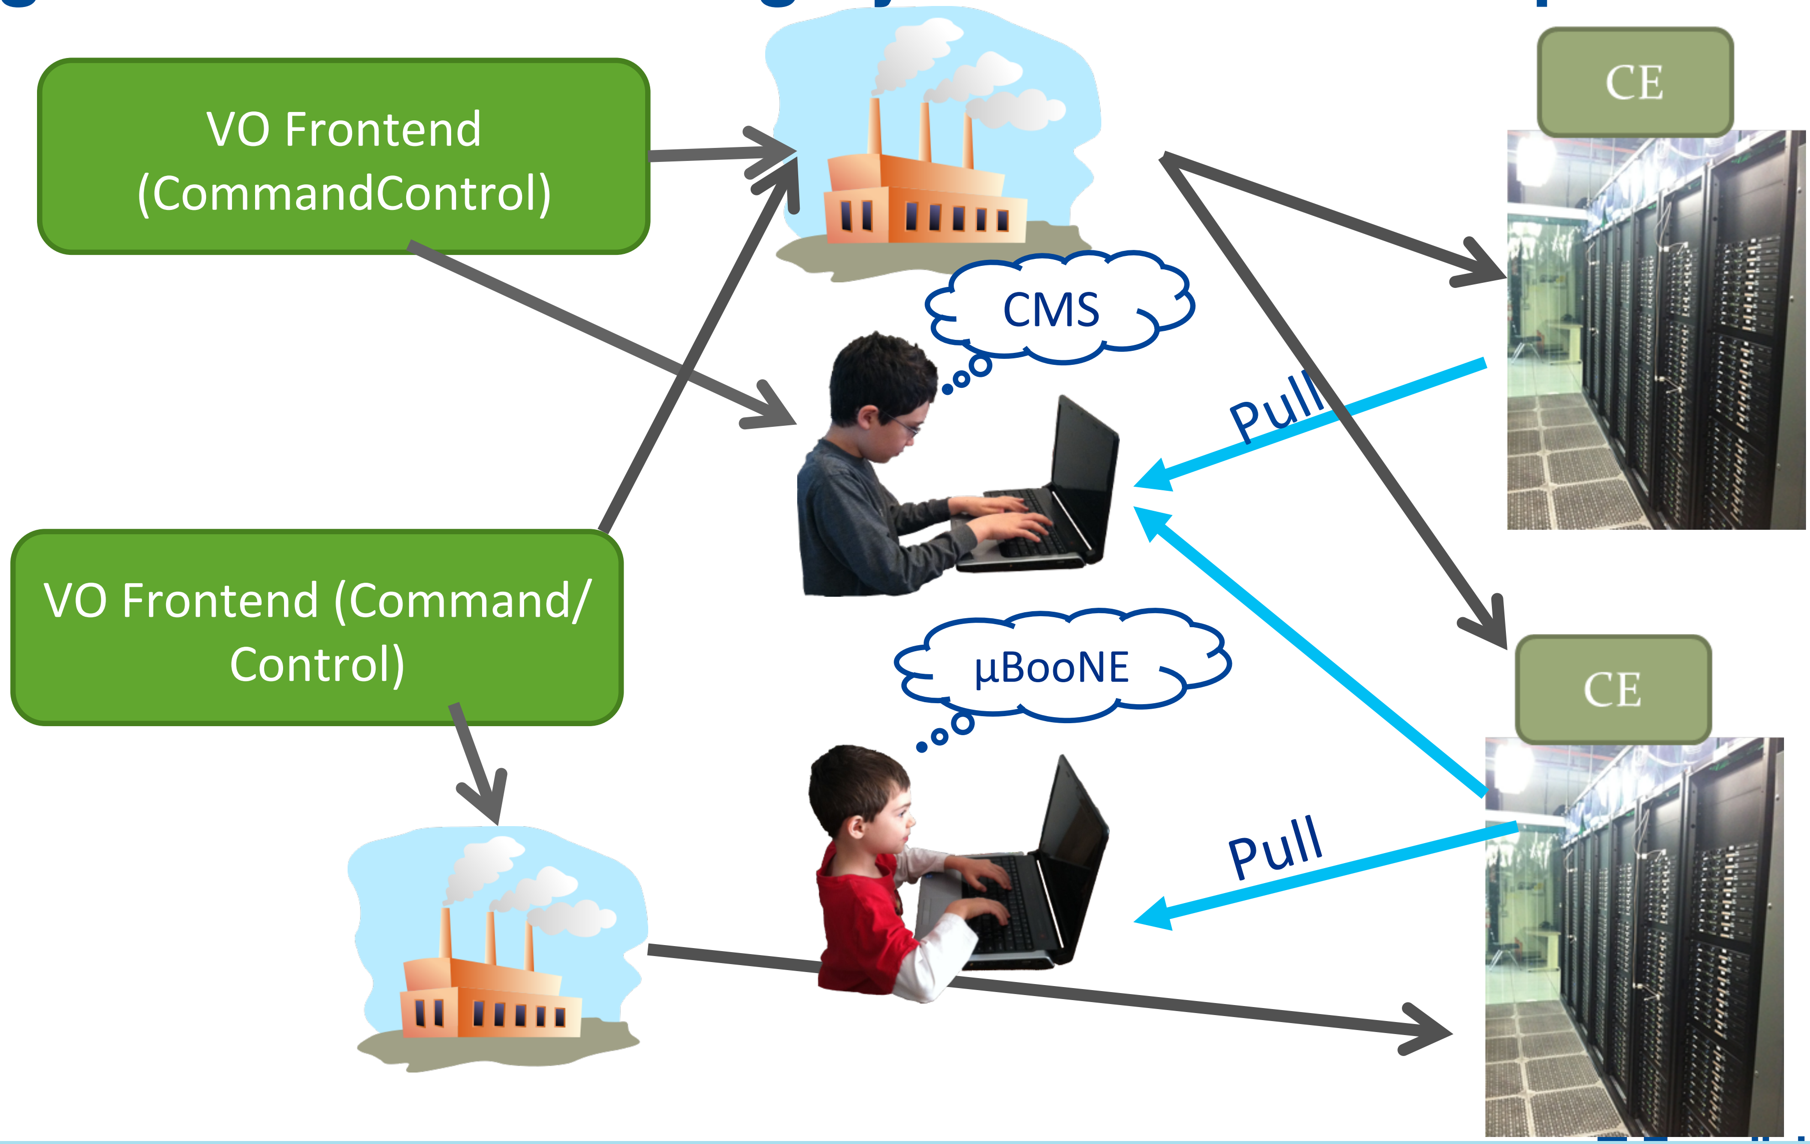

HEPCloud: Networking

- All models of distributing computing rely on the performance of the underlying (local and wide-area) network
- Fermilab is approaching **1 Terabit** data center – connect to Energy Sciences Network (ESNet) at 4*100 Gigabit
  - ESNet enables distributed computing beyond ESNet sites: 100 Gigabit peering points with other networks

- Zone-based security protection of network resources
- On-demand (**Software Defined Network**-based) traffic controls
- Virtualization of network resources

HEPCloud: Storage

- Data is the lifeblood of science
  - HEP experiments generate it by the station-wagon-load
  - Fermilab is a leader in the field in storing and serving petabytes of data to the world
- We are working with industry and other collaborators to modernize our services
  - Data storage and retrieval
  - Data cataloging
  - Support multiple-layer storage infrastructure approach
- One part of HEPCloud is to understand how to integrate all of these components – always driven by the experiment needs, both present and future