The Fermilab HEPCloud Facility: Adding 60,000 Cores for Science!

Burt Holzman, for the Fermilab HEPCloud Team

HTCondor Week 2016

May 19, 2016

My last Condor Week talk …

Condor in the CMS Experiment

Burt Holzman
Condor Week 2009
April 21, 2009

The Fermilab Facility: Today

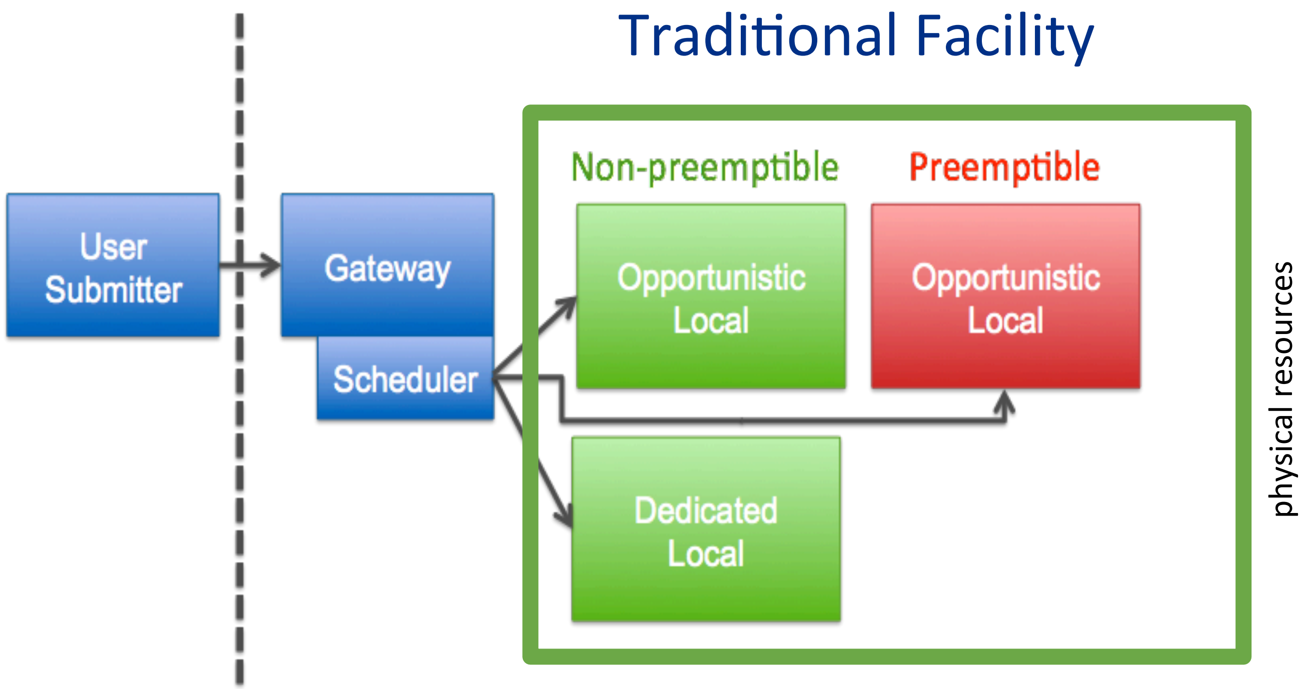

Drivers for Evolving the Facility

- HEP computing needs will be 10-100x current capacity
  - Two new programs coming online (DUNE, High-Luminosity LHC), while new physics search programs (Mu2e) will be operating

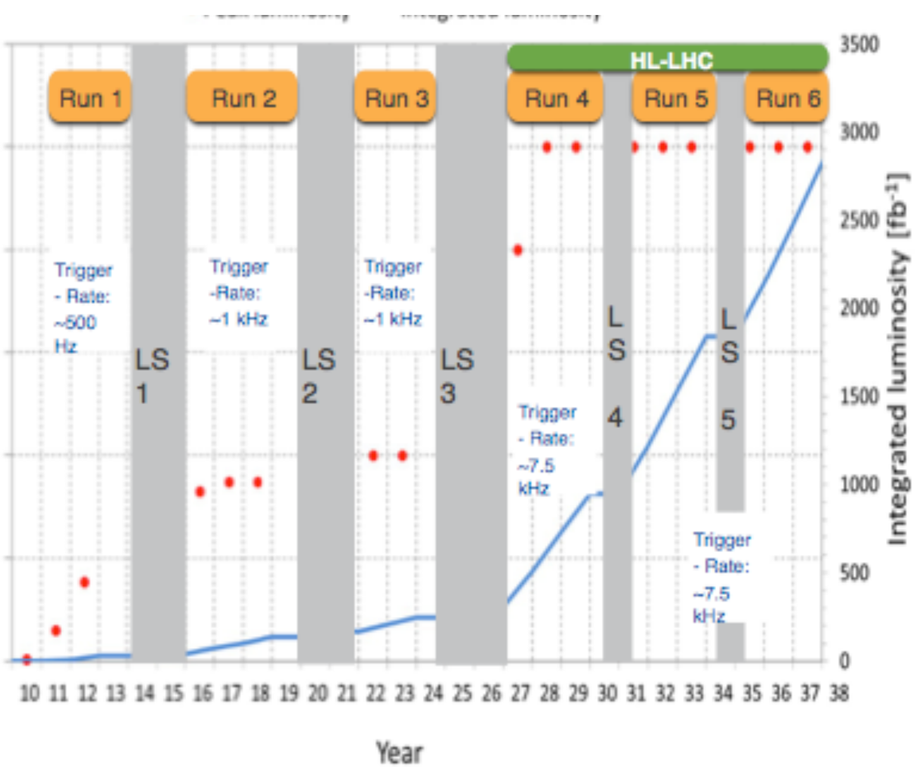

- Scale of industry at or above R&D
  - Commercial clouds offering increased **value** for decreased **cost** compared to the past

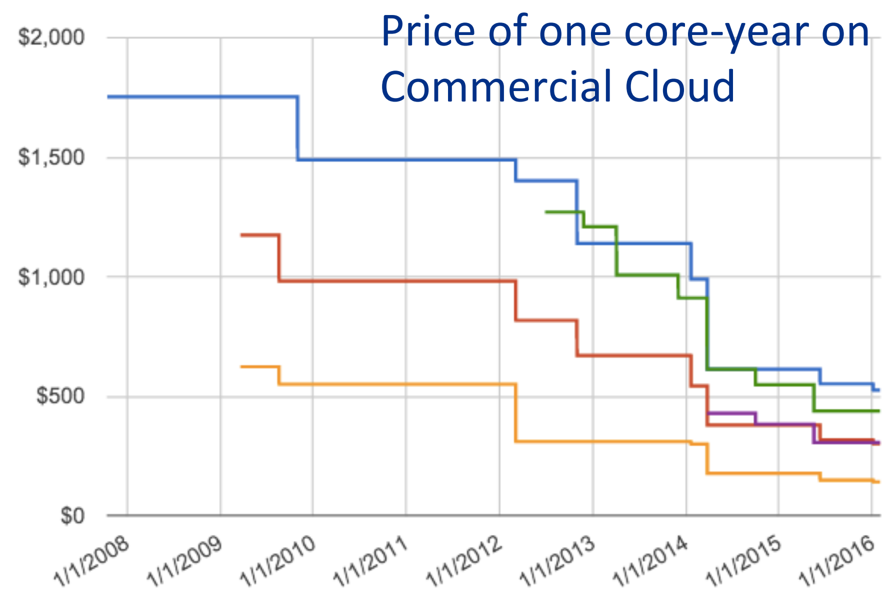

Drivers for Evolving the Facility: Elasticity

- Usage is not steady-state
- Computing schedules driven by real-world considerations (detector, accelerator, …) but also ingenuity – this is research and development of cutting-edge science

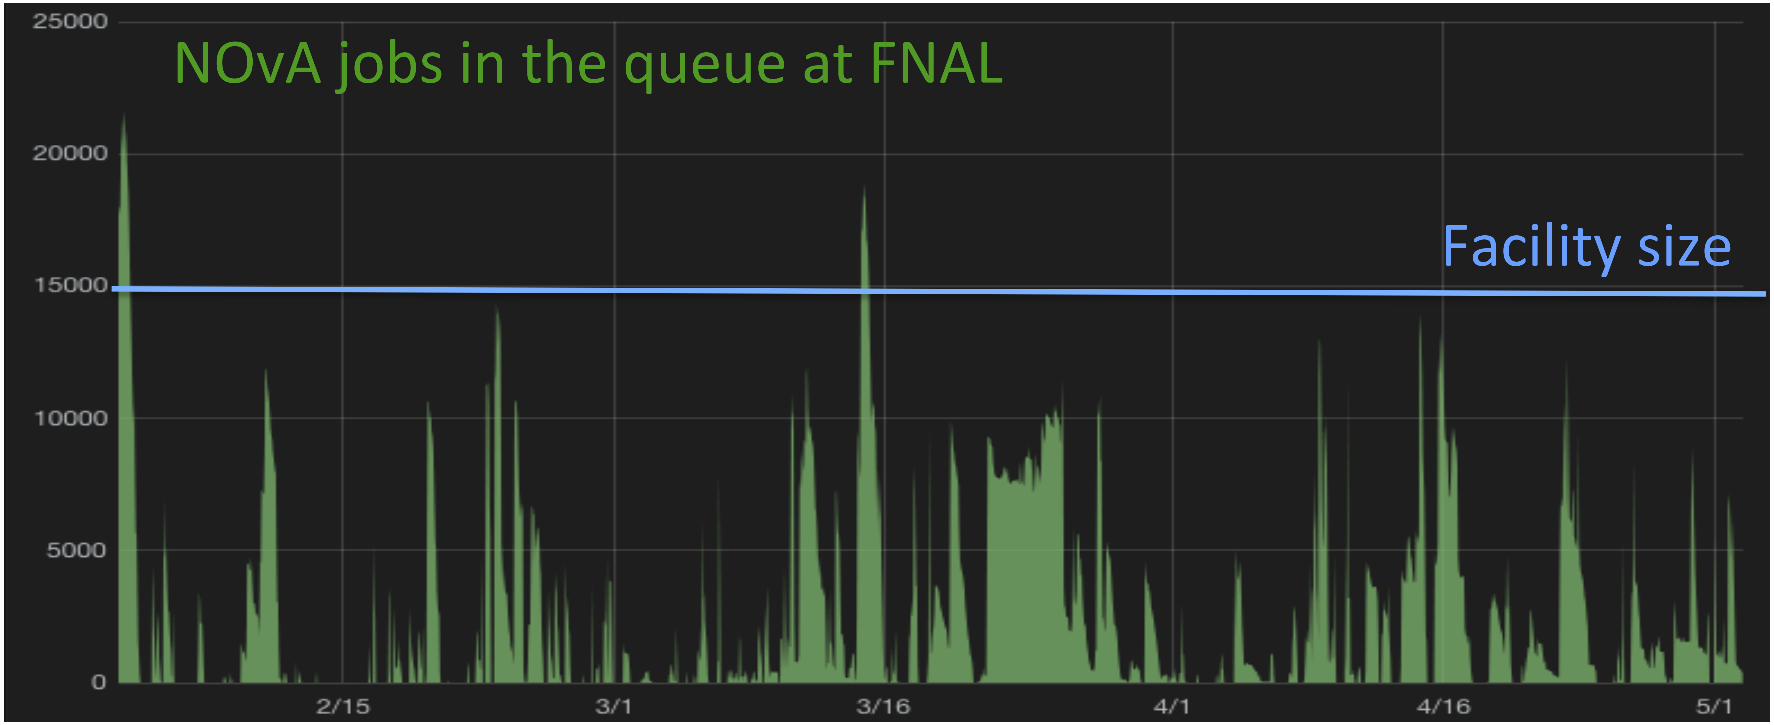

Classes of Resource Providers

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

The Fermilab Facility: Today

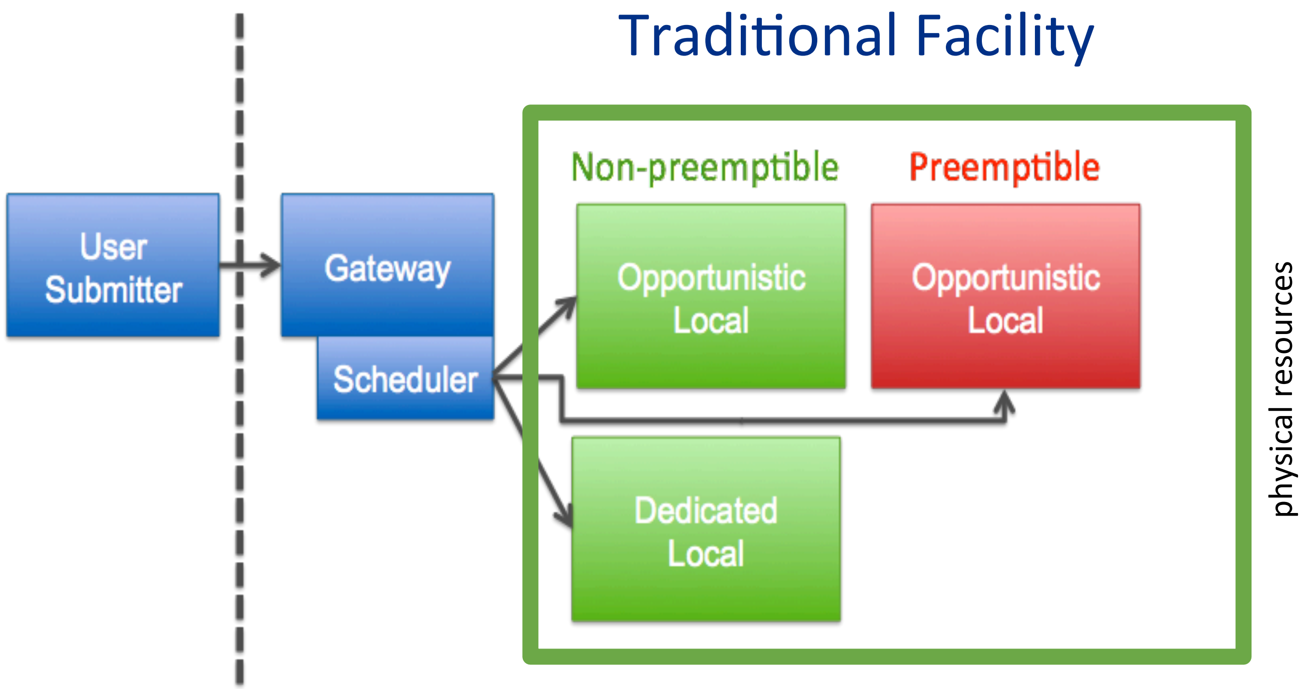

The Fermilab Facility: ++Today

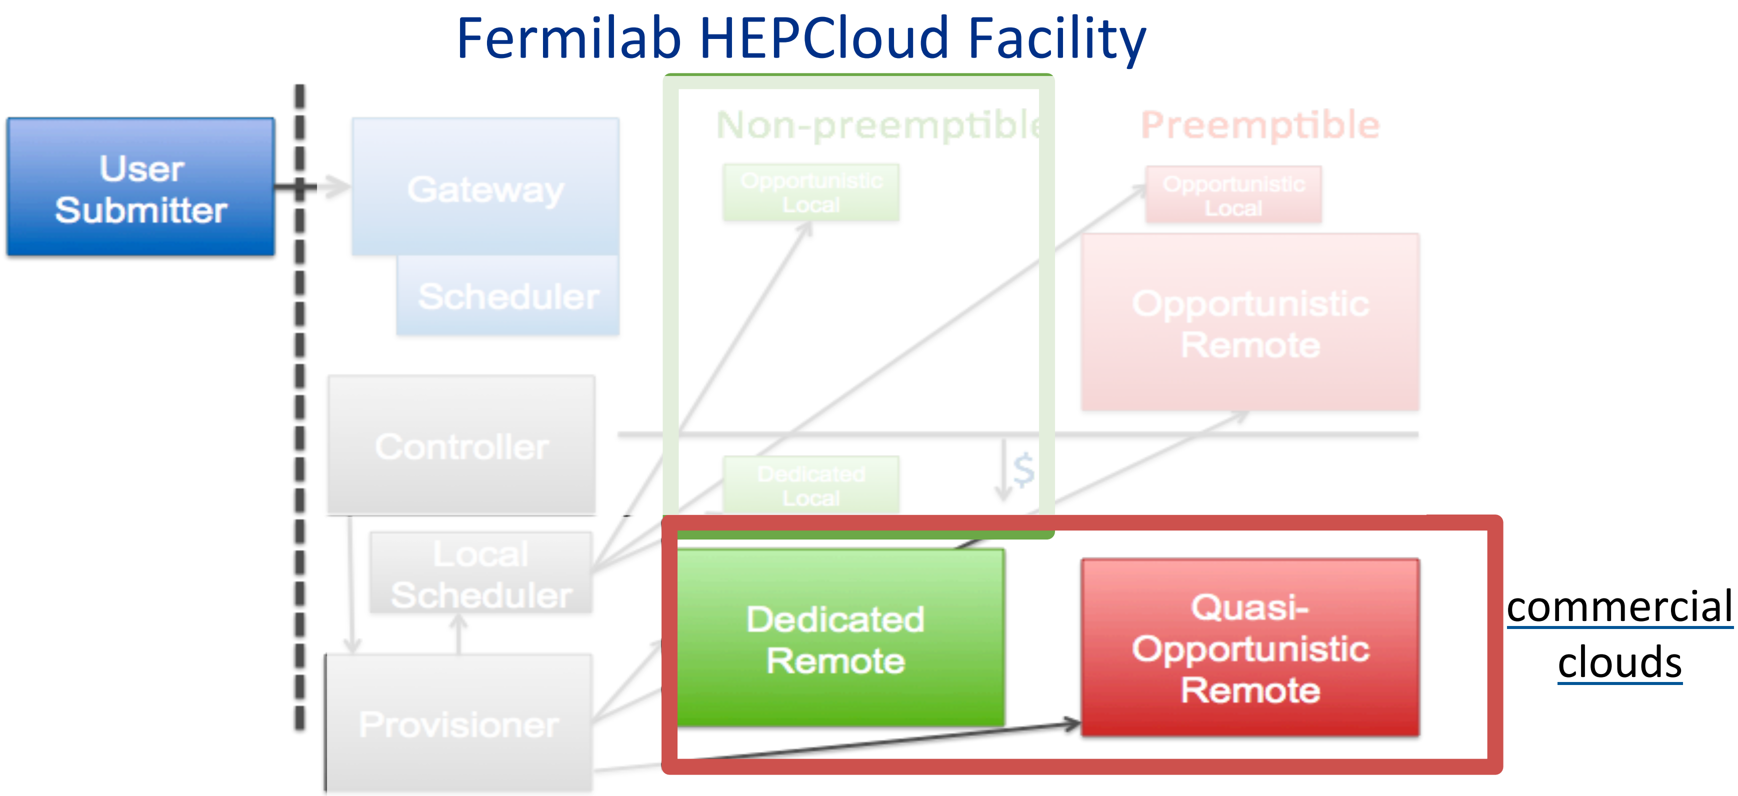

- Provision commercial cloud resources in addition to physically owned resources
- Transparent to the user
- Pilot project / R&D phase

HEPCloud Collaborations

- Engage in collaboration to leverage tools and experience whenever possible

- **HTCondor** – common provisioning interface
  - Foundation underneath **glideinWMS**
  - Grid technologies – Open Science Grid, Worldwide LHC Computing Grid
  - Preparing communities for distributed computing
- CMS – collaborative knowledge and tools, cloud-capable workflows
- BNL and ATLAS – engaged in next HEPCloud phase

HEPCloud Collaborations

- Engage in collaboration to leverage tools and experience
  w
- H
  -
- G ... C
  C
  -
- C ... e
  w
- B

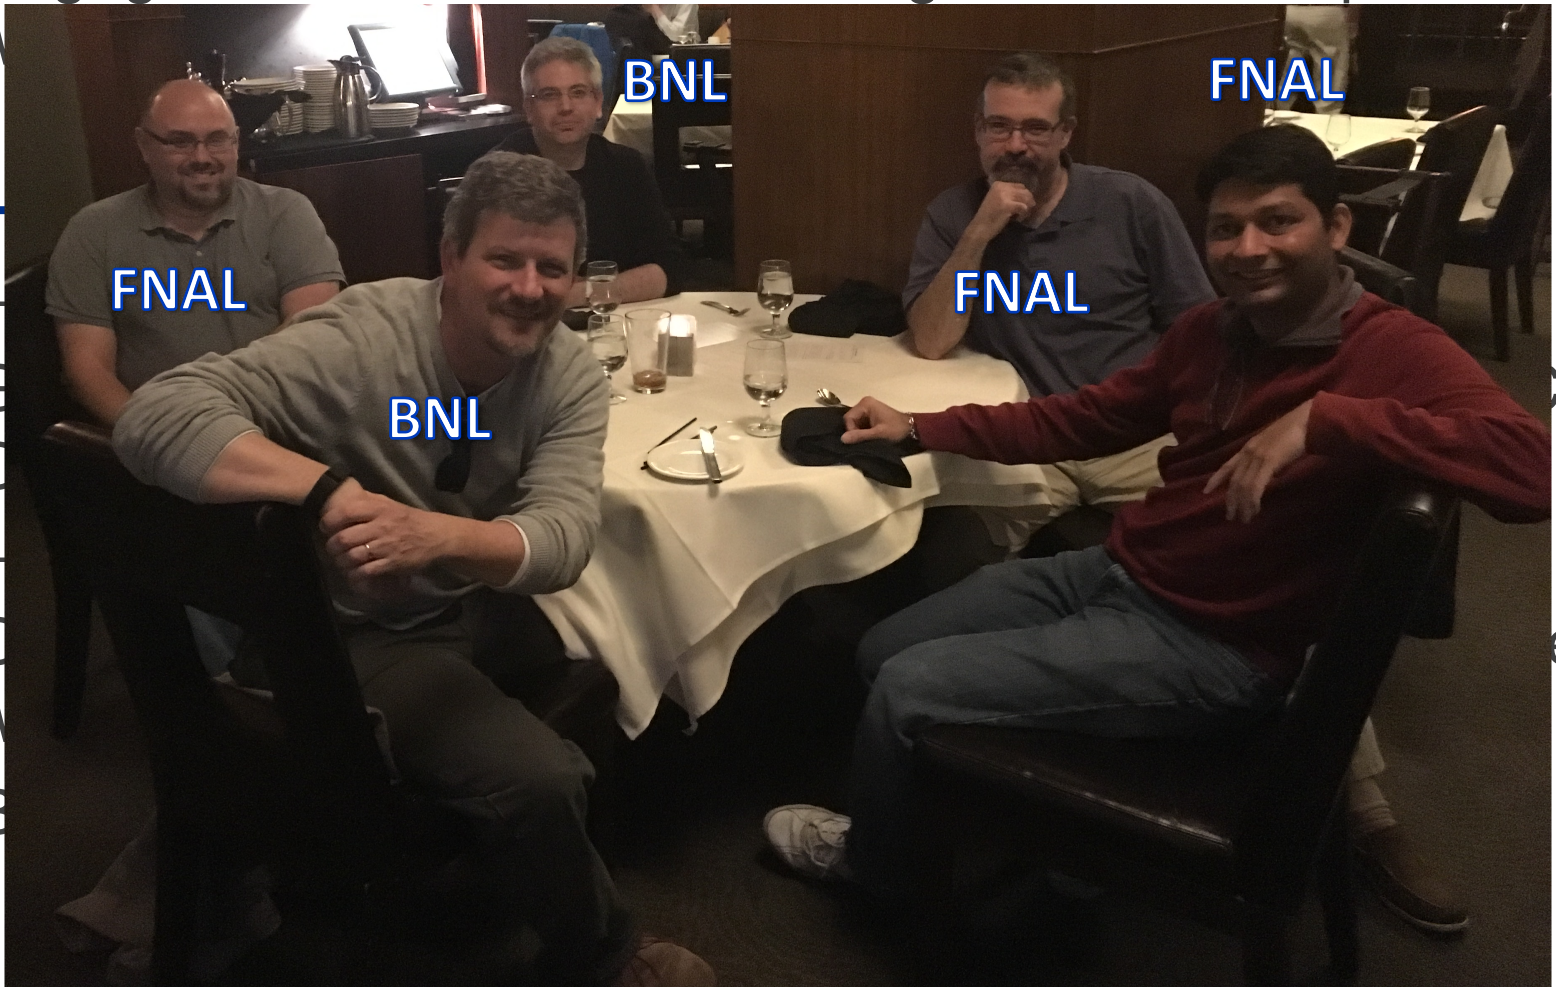

Fermilab HEPCloud: expanding to the Cloud

- Where to start?
  - Market leader: **Amazon Web Services (AWS)**

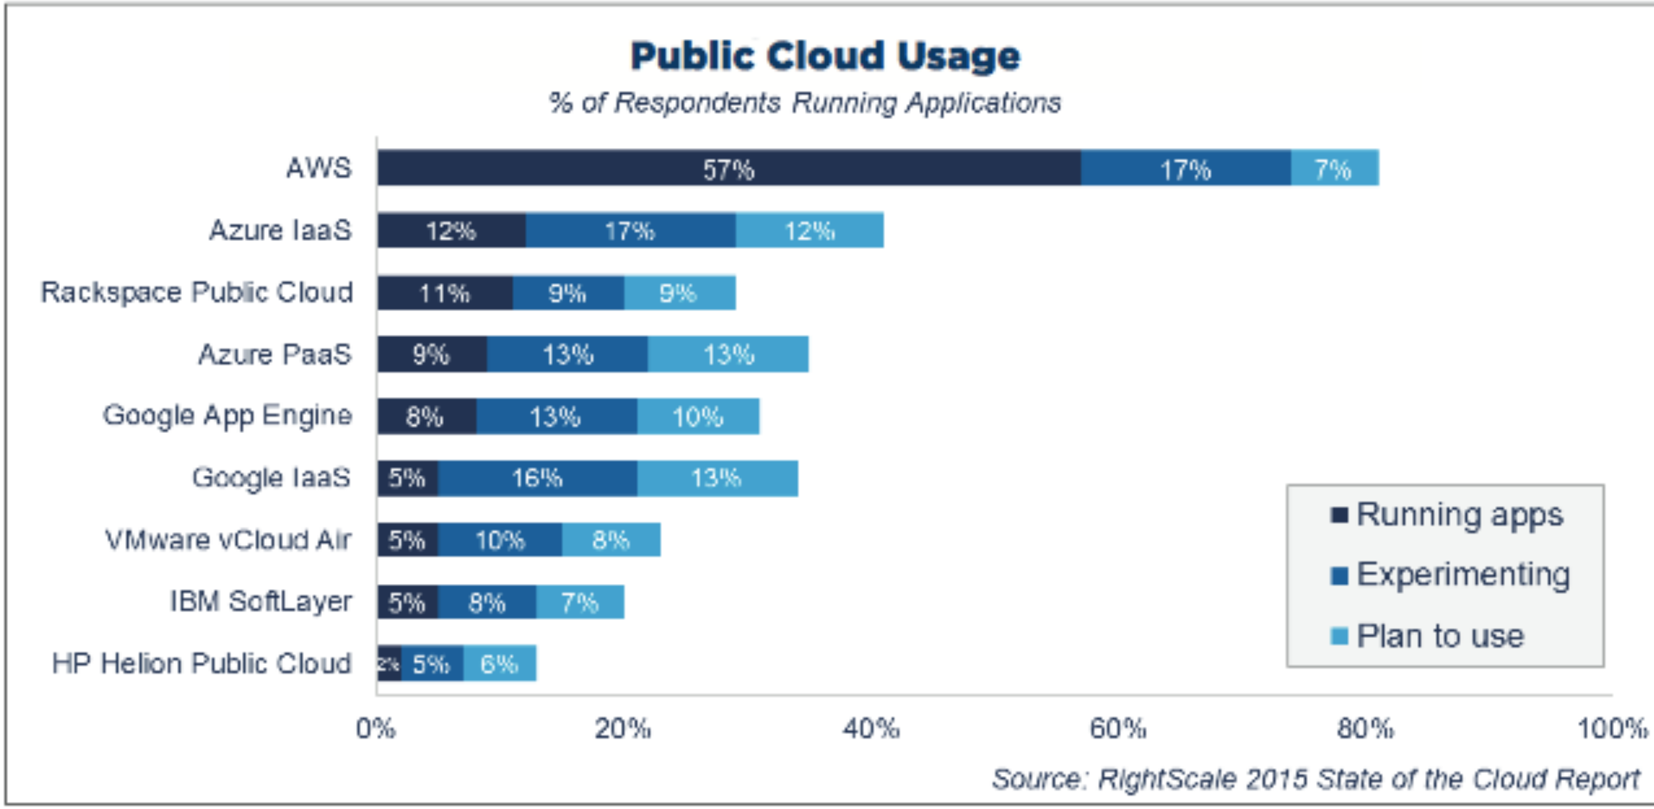

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof.

AWS topology – three US data centers ("regions")

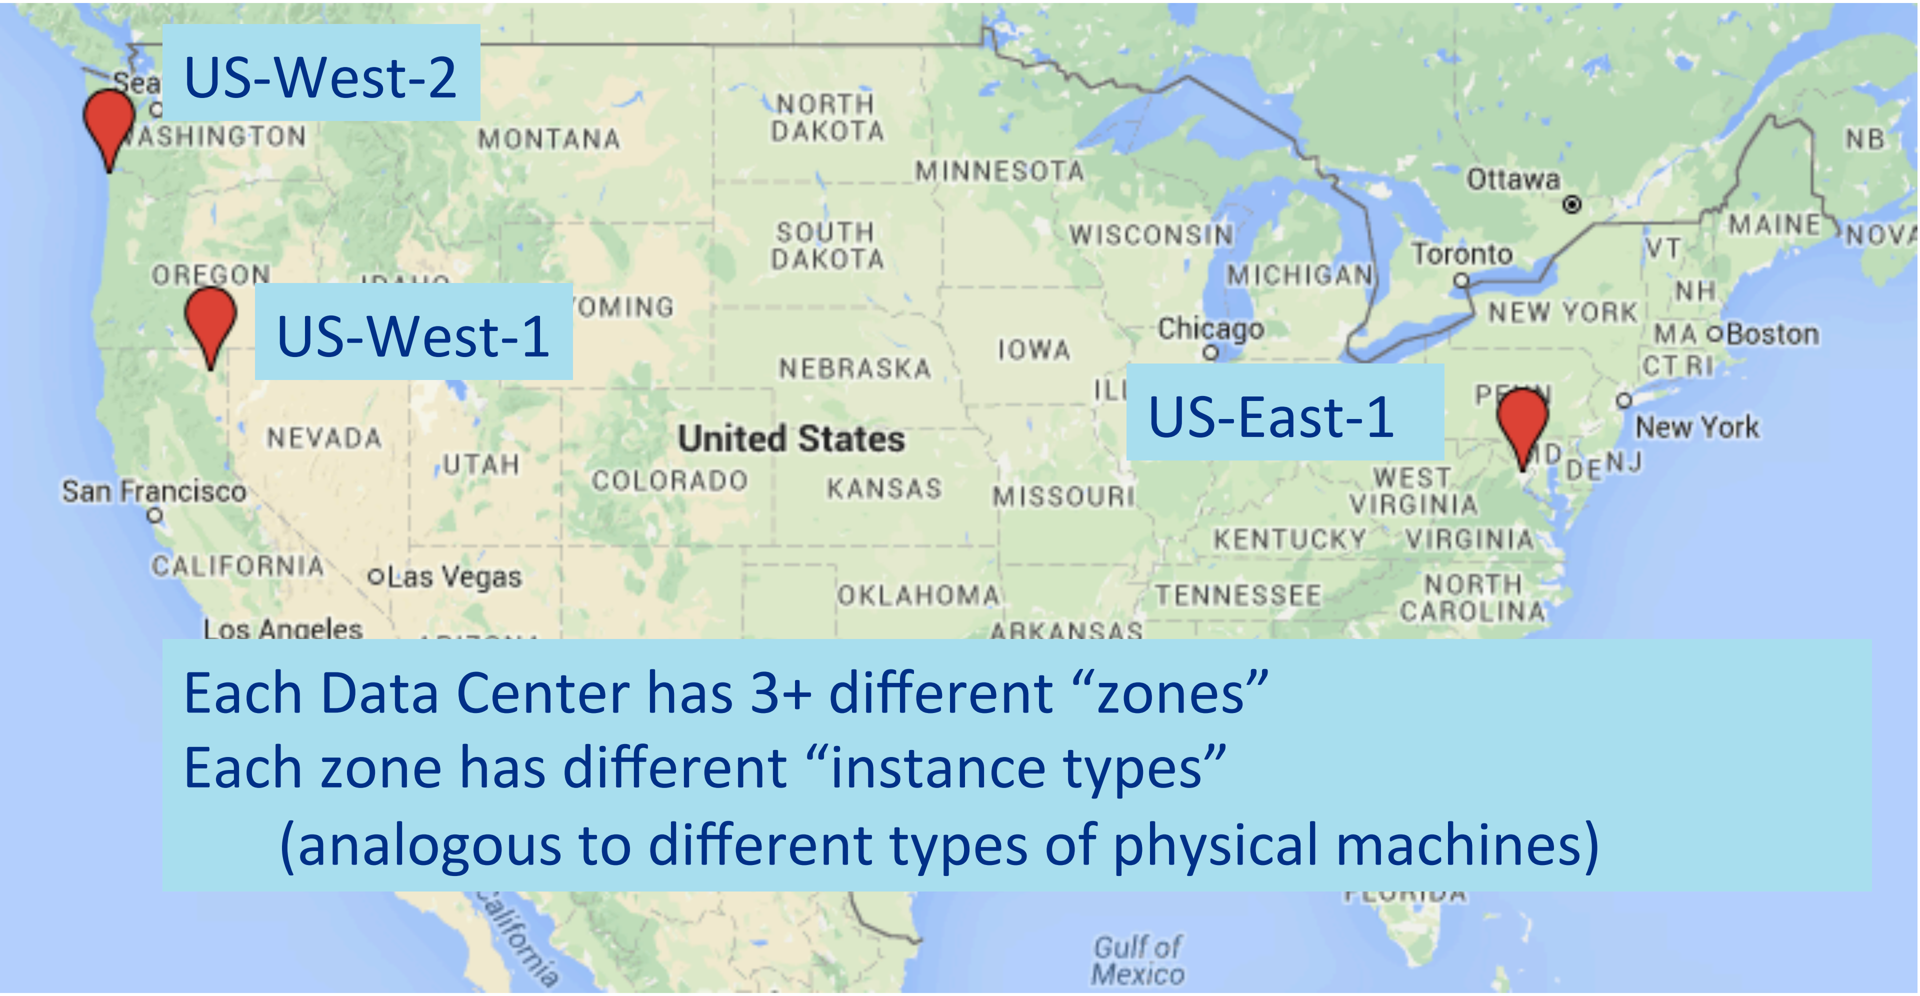

Each Data Center has 3+ different "zones"
Each zone has different "instance types"
(analogous to different types of physical machines)

Pricing: using the AWS “Spot Market”

- 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
  - If (market price < bid), you pay only market price for the provisioned resource
    - If (market price > bid), you don’t get the resource
  - If the price fluctuates while you are running and the market price exceeds your original bid price, you may get kicked off the node (with a 2 minute warning!)

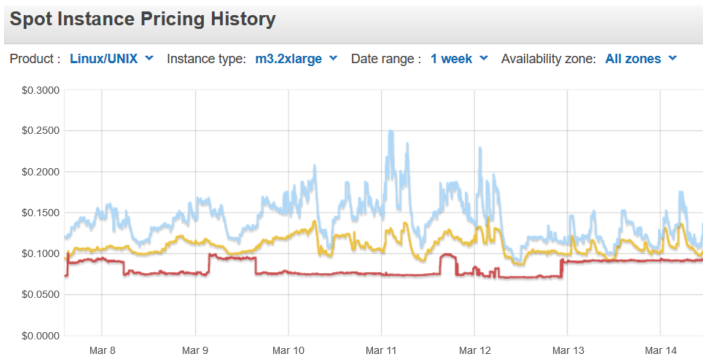

Some HEPCloud Use Cases

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

NOvA: Neutrino Experiment

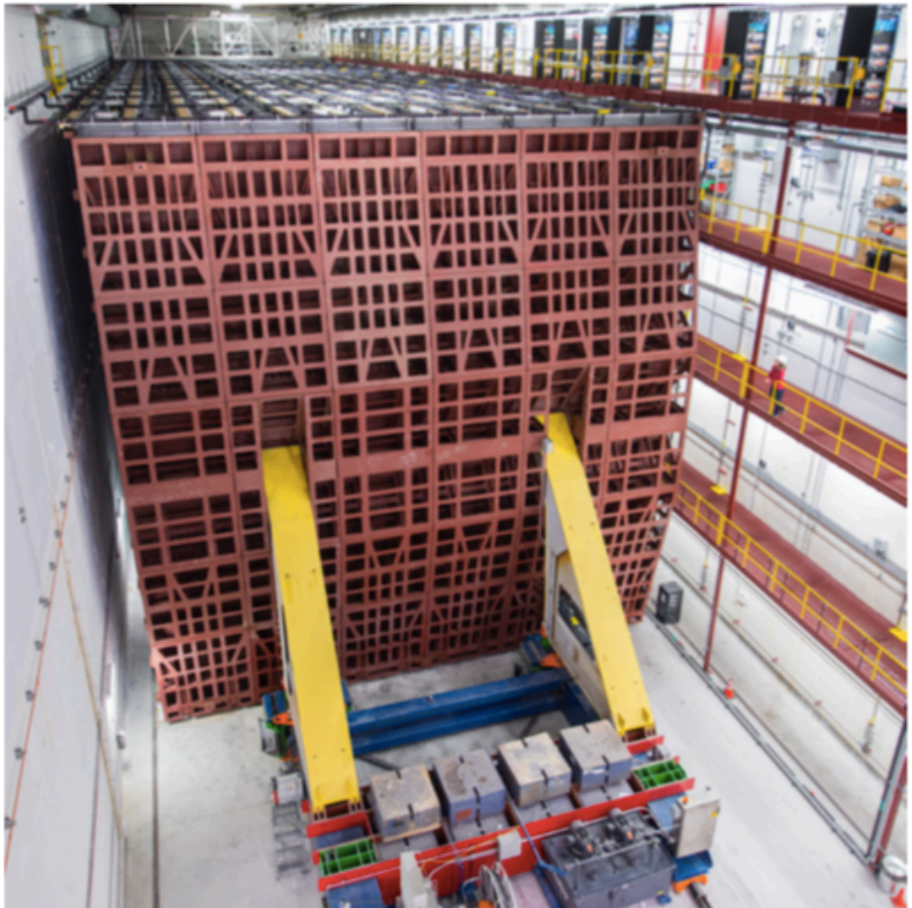

The NOvA detector in Minnesota occupies an area about the size of two basketball courts. It is 200 feet long and made of modules 50 feet high and 50 feet wide. The detector records particle tracks from neutrinos sent by a powerful accelerator at Fermilab. The construction of the NOvA detectors was completed 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

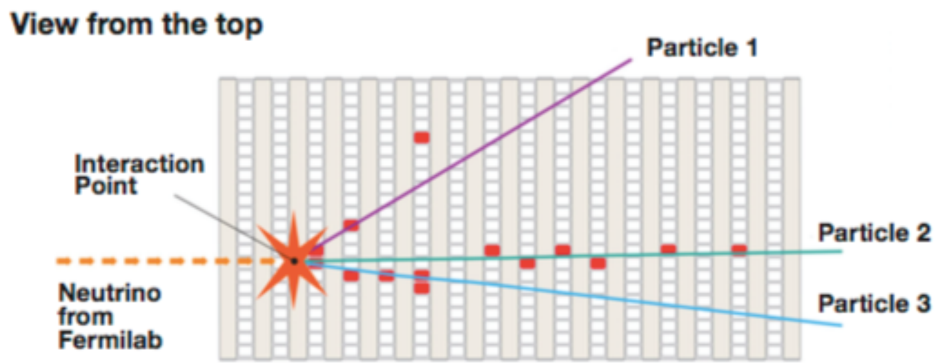

Neutrinos rarely interact with matter. When a neutrino smashes into an atom in the NOvA detector in Minnesota, it creates distinctive particle tracks. Scientists explore these particle interactions to better understand the transition of muon neutrinos into electron neutrinos. The experiment also helps 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

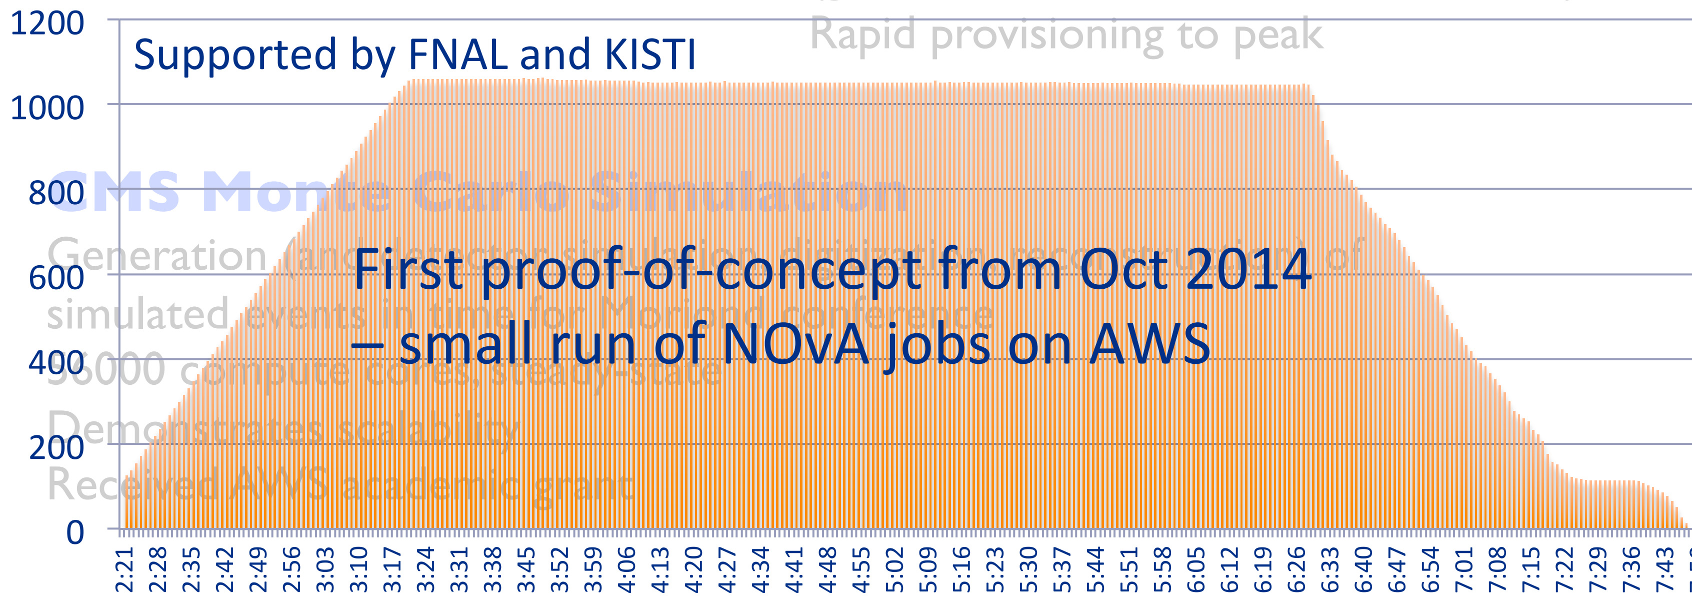

NOvA Use Case – running at 4k cores

- Added support for general-purpose data-handling tools (SAM, IFDH, F-FTS) for AWS Storage and used them to stage both input datasets and job outputs

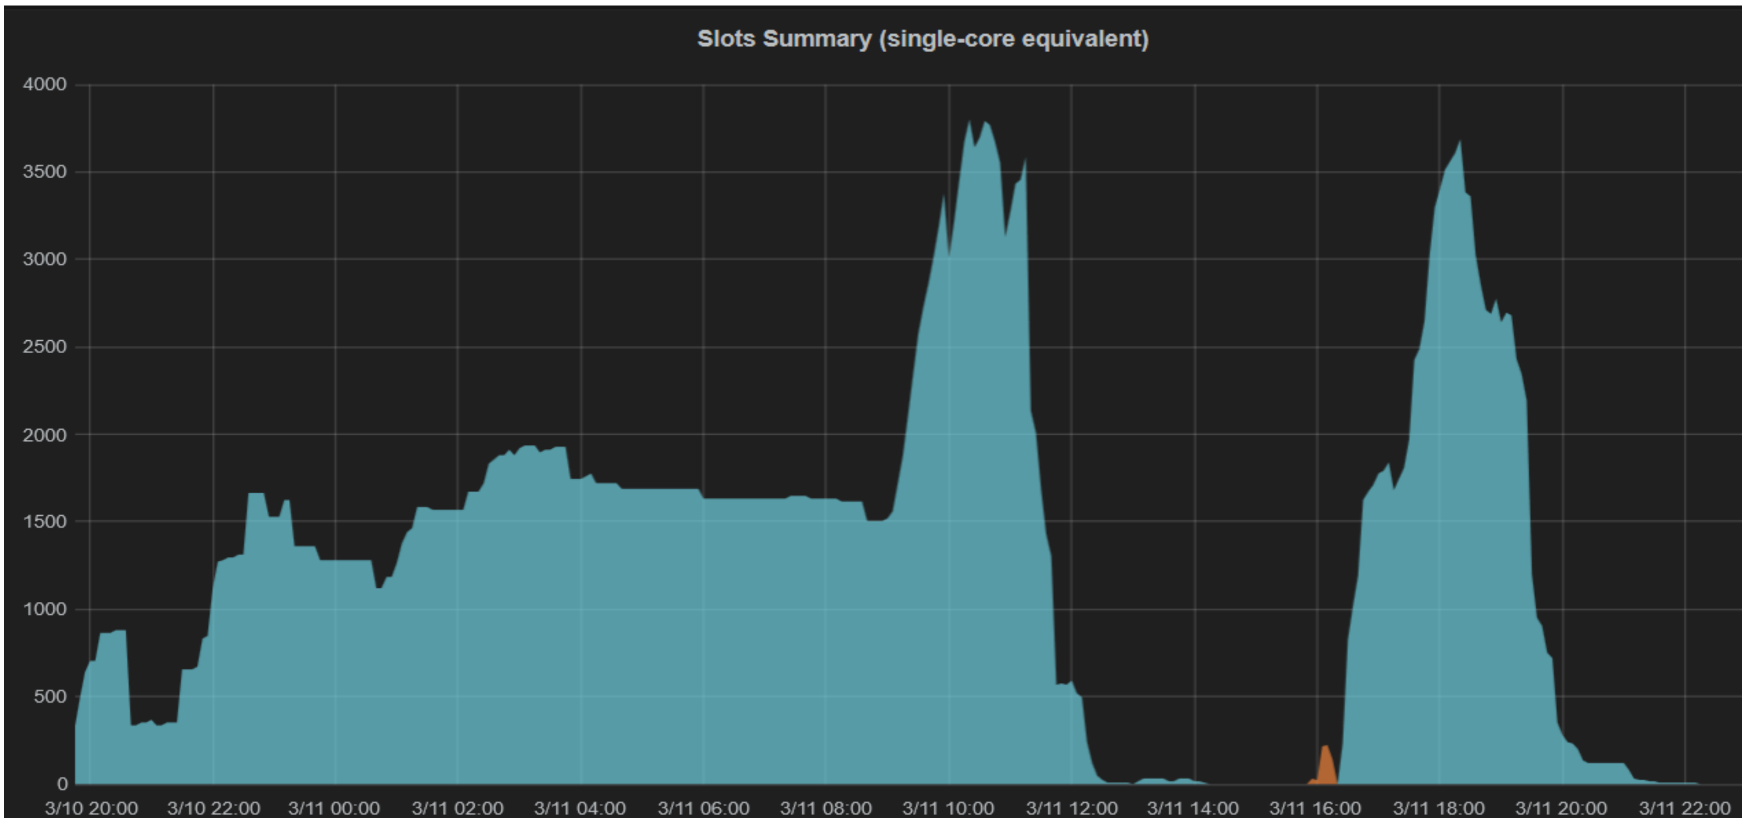

Some HEPCloud Use Cases

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

CMS: Large Hadron Collider Experiemnt

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

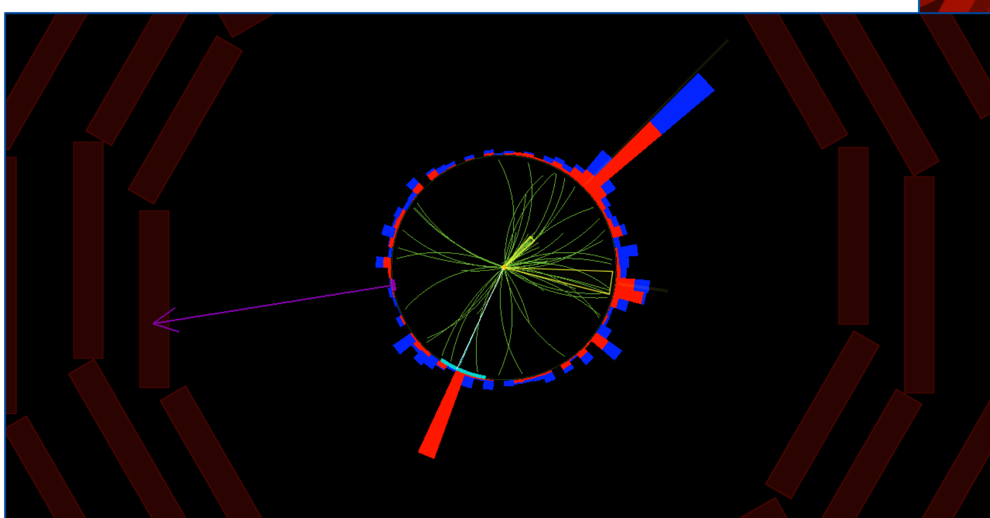

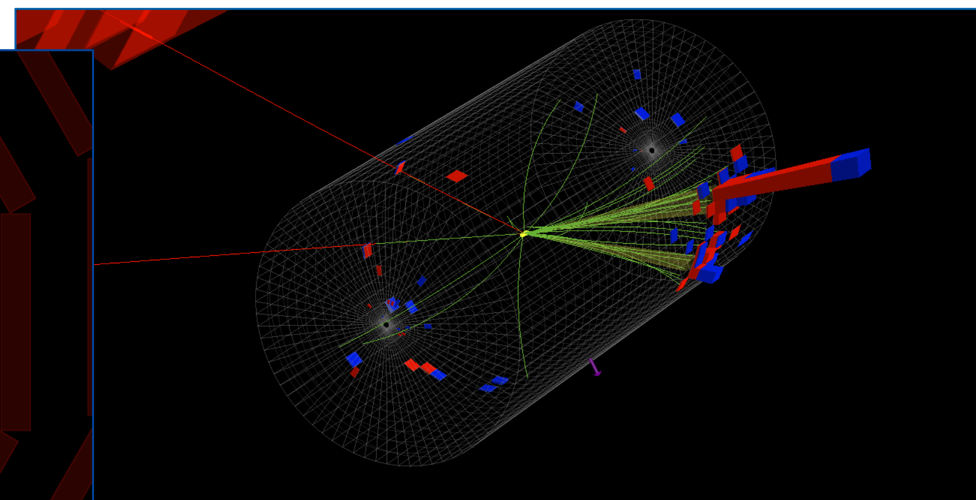

Reaching ~60k slots on AWS with FNAL HEPCloud

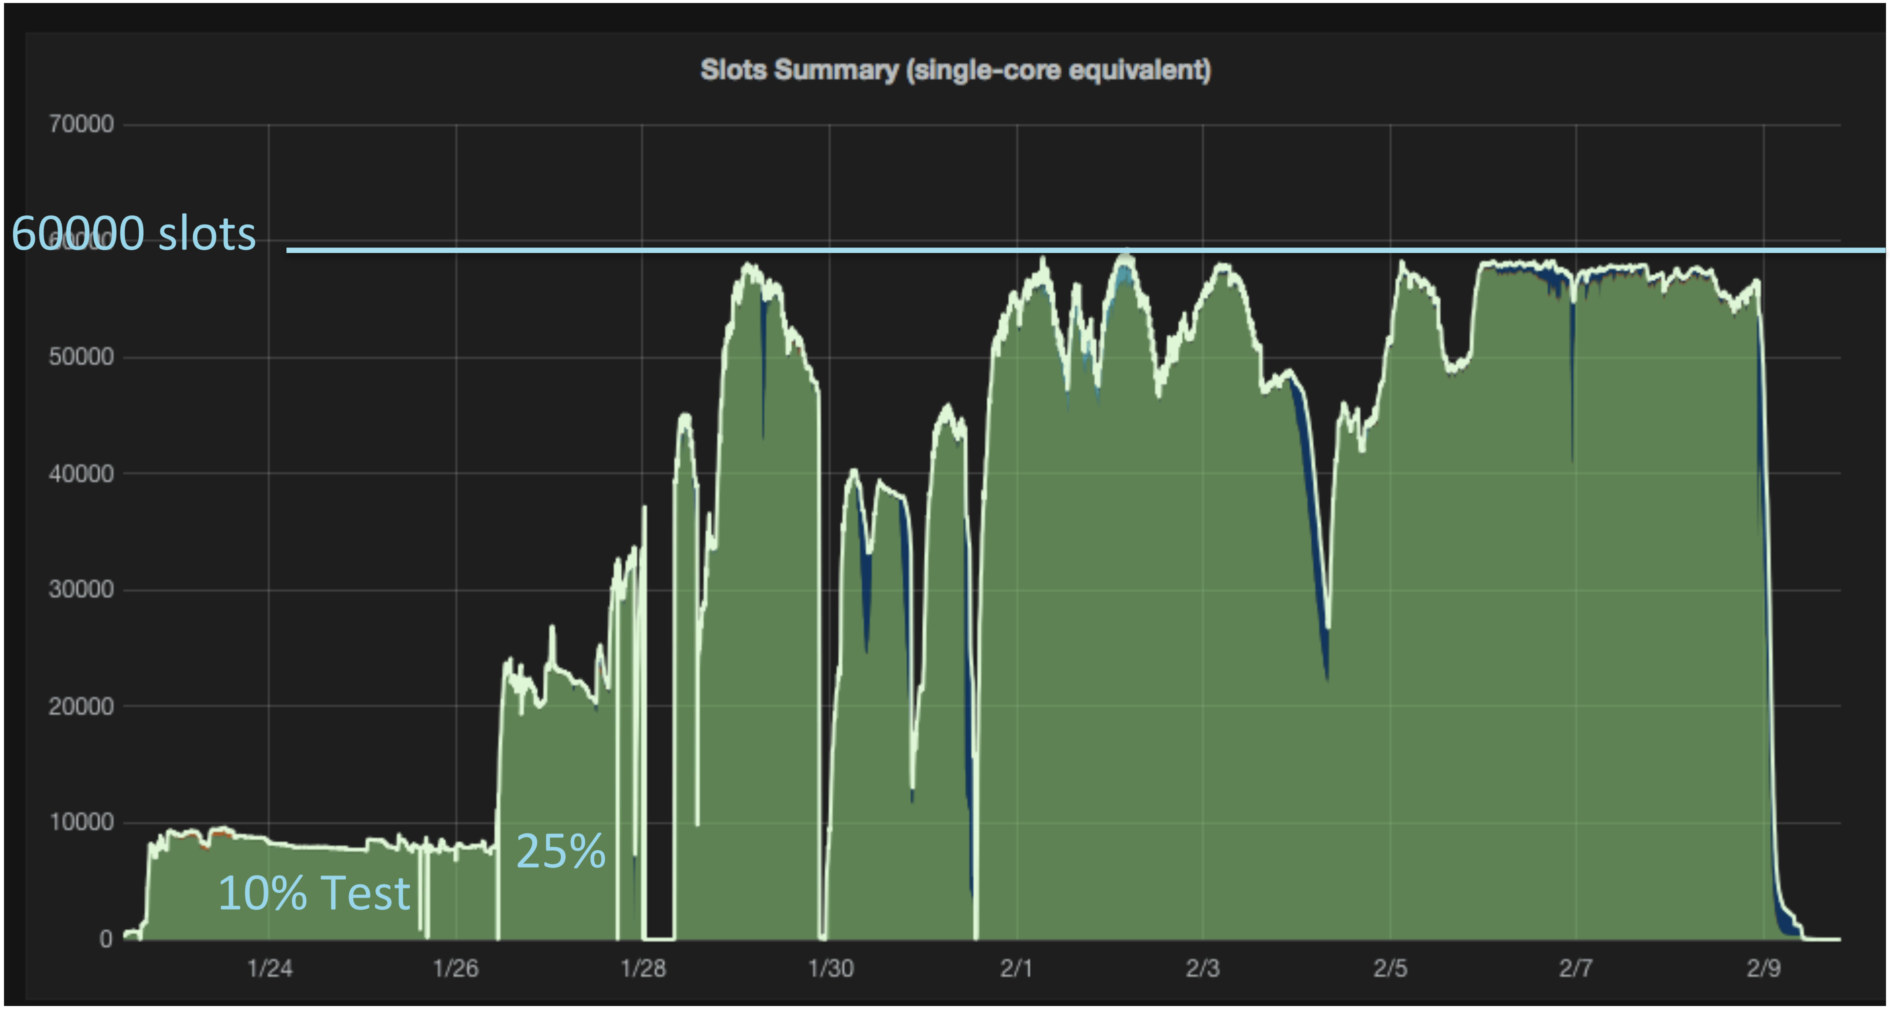

HEPCloud AWS slots by Region/Zone

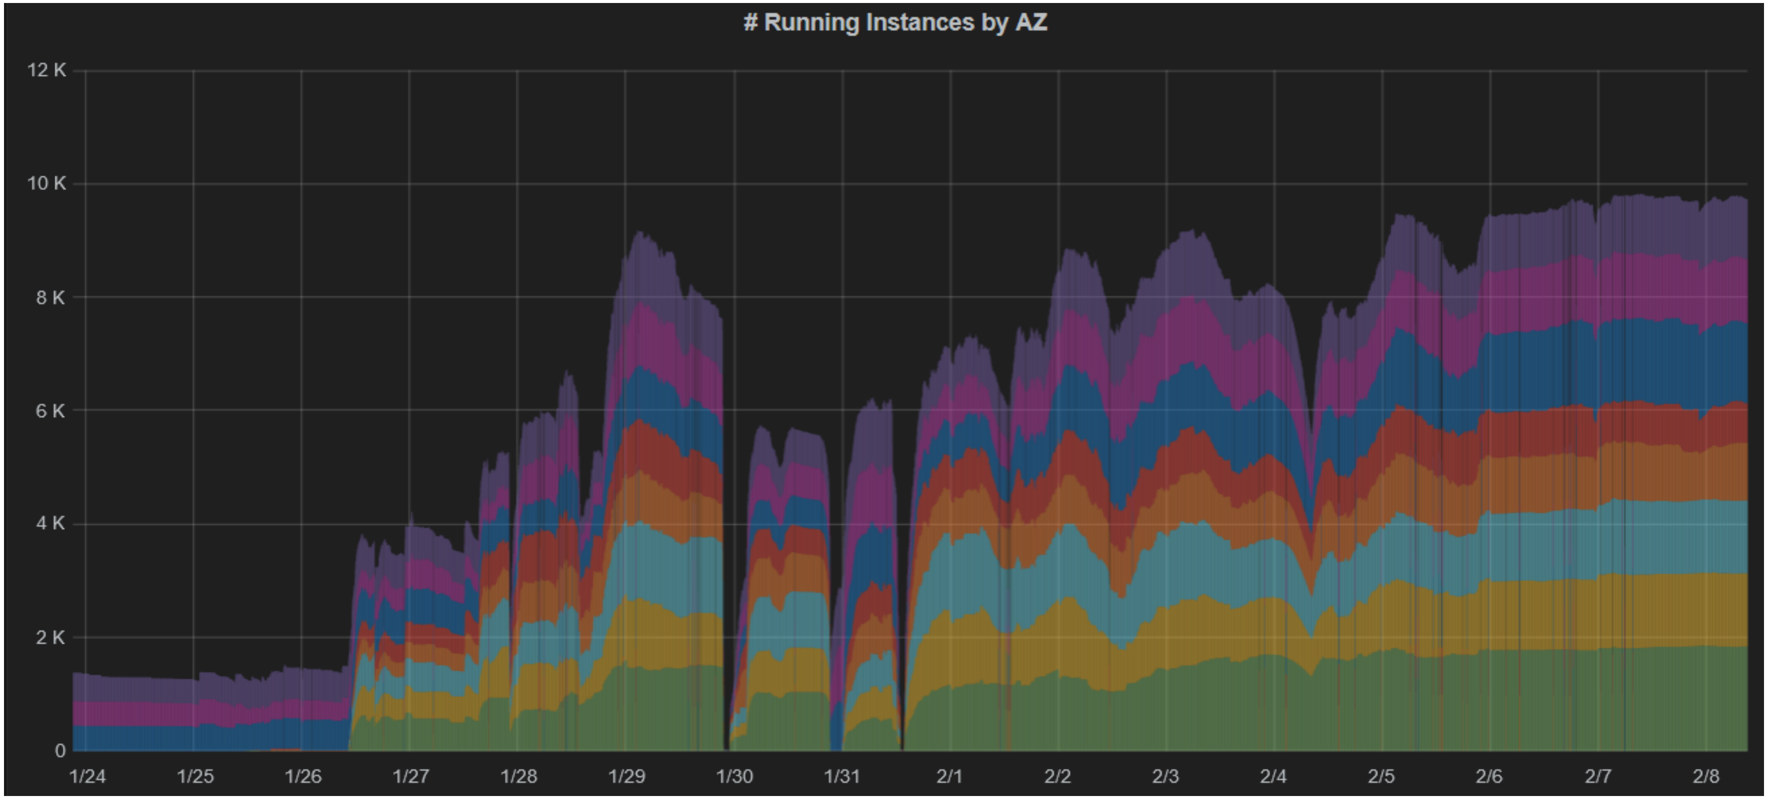

Each color corresponds to a different region+zone

HEPCloud AWS slots by Region/Zone/Type

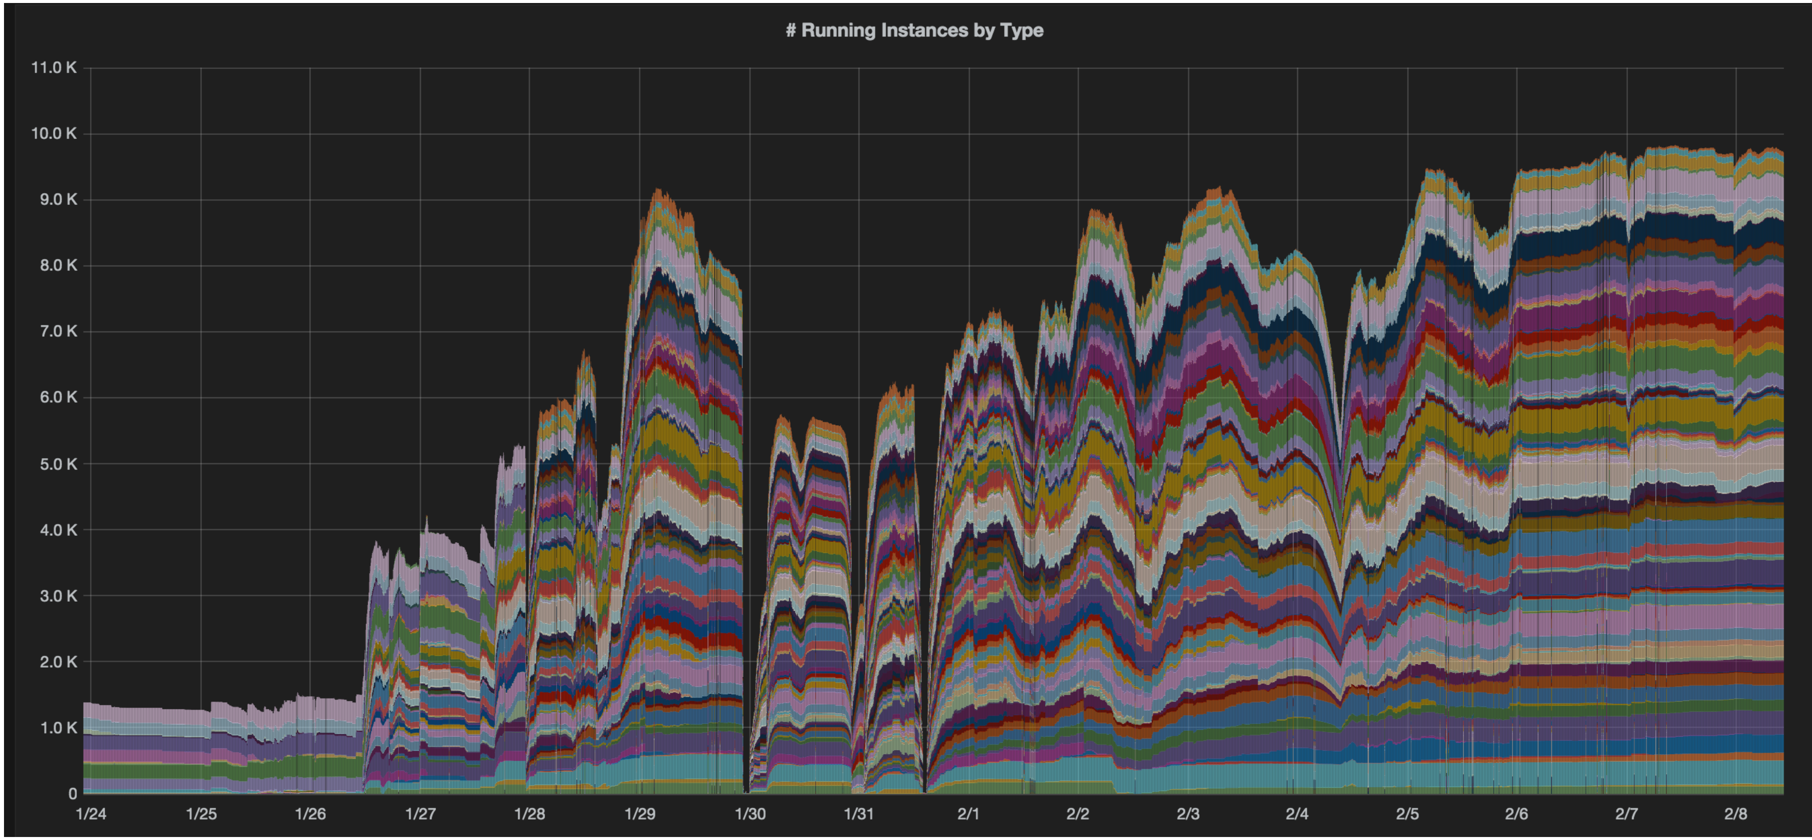

Each color corresponds to a different region+zone+machine type

HEPCloud/AWS: 25% of CMS global capacity

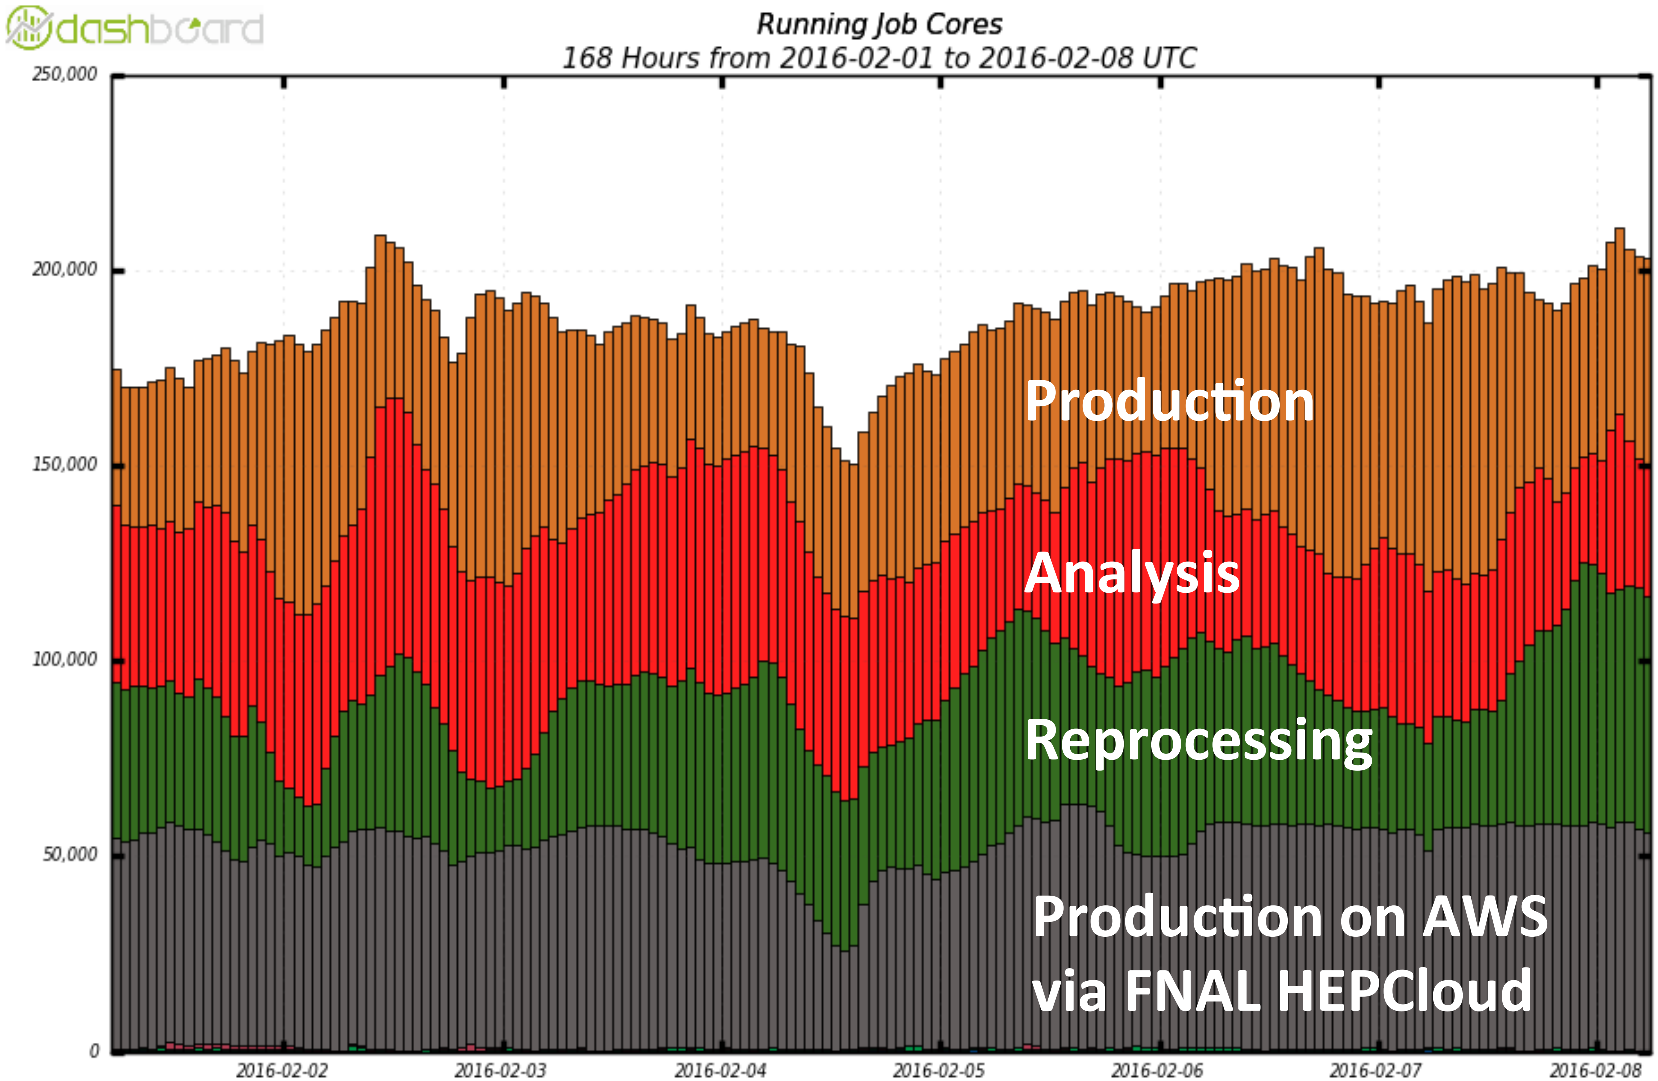

Fermilab HEPCloud compared to global CMS Tier-1

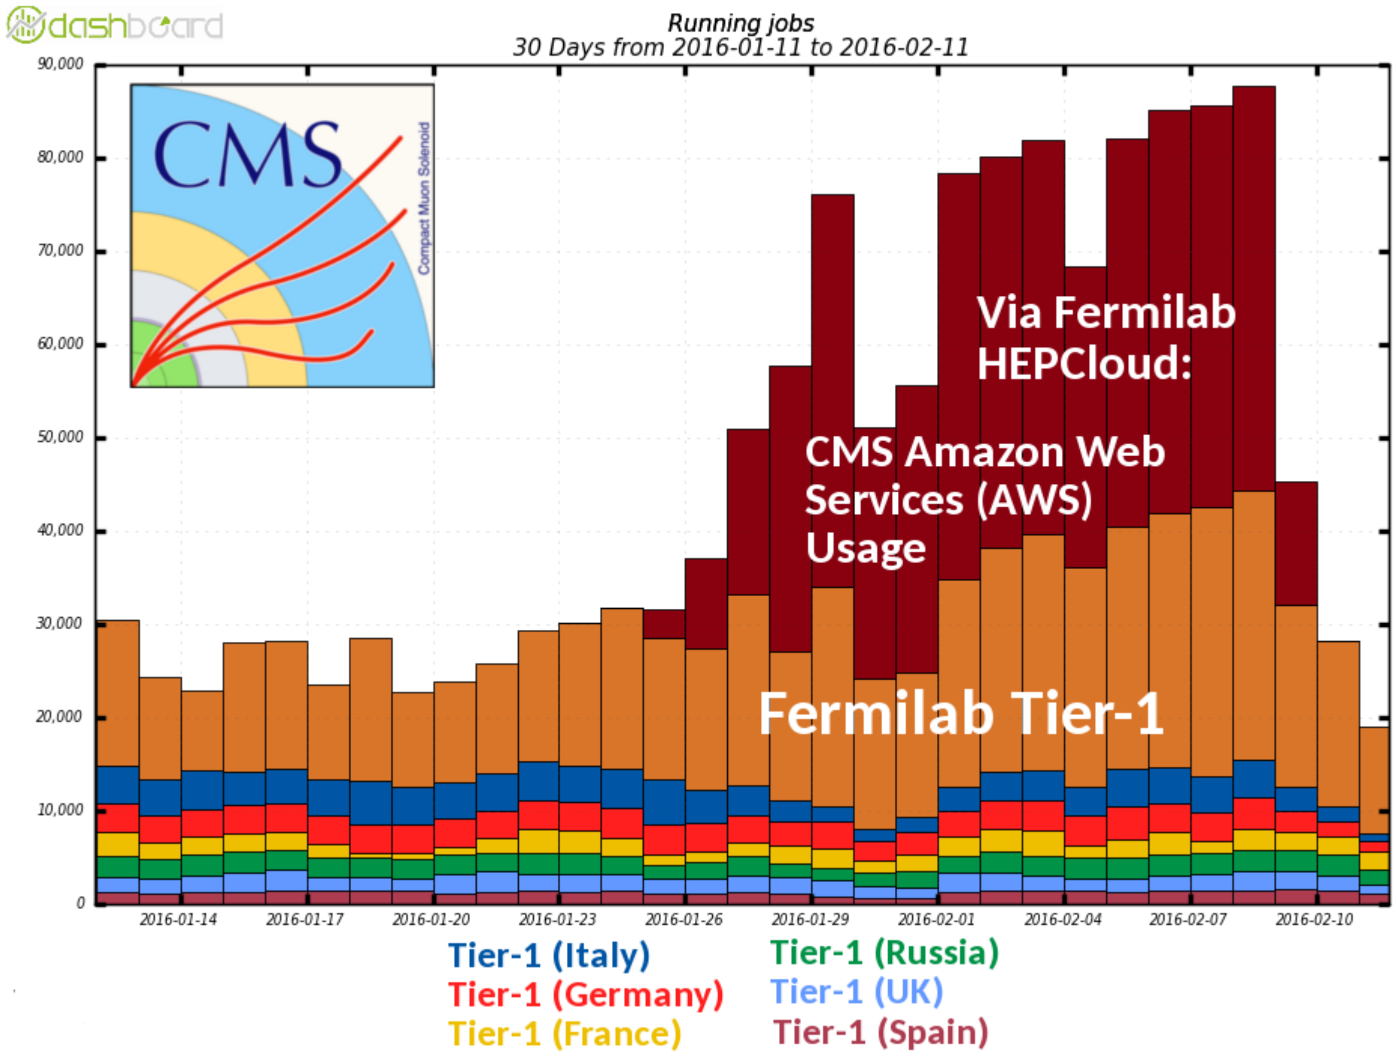

HEPCloud: Orchestration

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

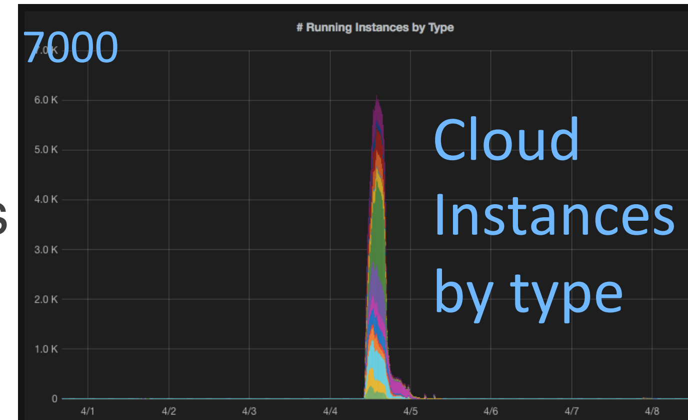

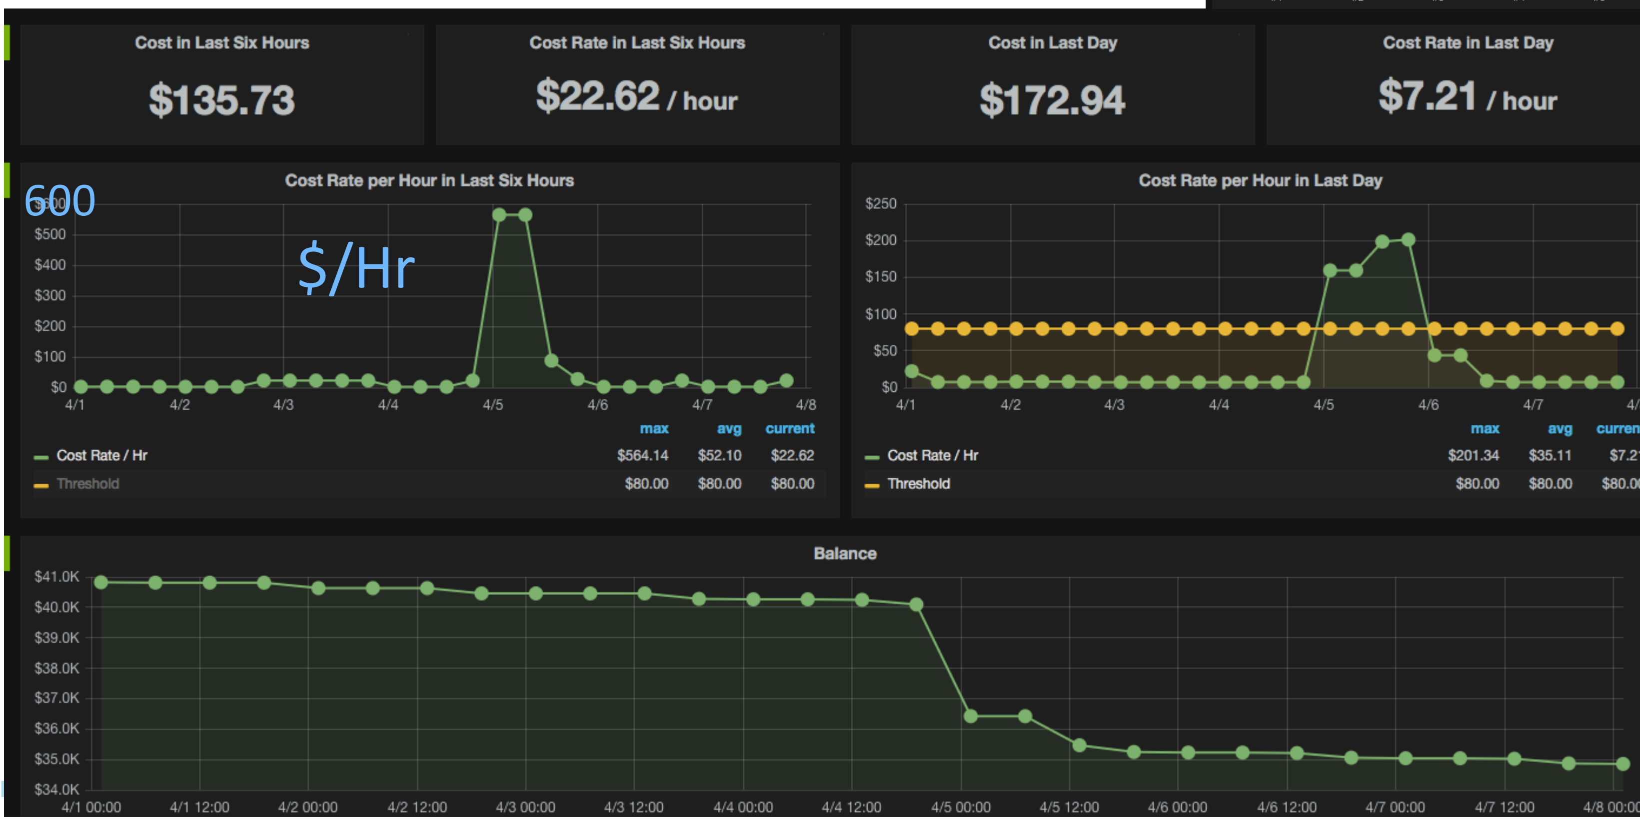

On-premises vs. cloud cost comparison

- Average cost per core-hour
  - On-premises resource: **.9** cents per core-hour
    - Includes power, cooling, staff
  - Off-premises at AWS: **1.4** cents per core-hour
    - Ranged up to 3 cents per core-hour at smaller scale
- Benchmarks
  - Specialized ("ttbar") benchmark focused on HEP workflows
    - On-premises: **0.0163** (higher = better)
    - Off-premises: **0.0158**

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

HTCondor: Critical to Success

- **All resources** provisioned with HTCondor
- **First test** of EC2 GAHP at scale
  - Worked* with HTCondor team to improve EC2 GAHP
    - Improved stability of GAHP (less mallocs)
    - Improved Gridmanager response to crashed GAHP
    - Reduce number of EC2 API calls and exponential backoff (BNL request)
- We need a agent to speak to **bulk provisioning APIs**
- **condor_annex** (see next talk)
  - We want HTCondor to provision the "big three"
    - Amazon EC2
    - Google Cloud Engine
    - Microsoft Azure
  - condor_annex should be part of the HTCondor ecosystem (ClassAds, integration with condor tools, run as non-privileged user)

* Todd M codes, we test

Thanks

- HTCondor team
- CMS and NOvA experiments
- The glideinWMS project
- FNAL HEPCloud Leadership Team: Stu Fuess, Gabriele Garzoglio, Rob Kennedy, Steve Timm, Anthony Tiradani
- Open Science Grid
- Energy Sciences Network
- Amazon Web Services
- ATLAS/BNL for initiating work with AWS team (and for providing some introduction in John Hover's talk yesterday!)