



Cloud Resources and the OneFacility



RECENT CLOUD WORK AND FUTURE PLANS

**STEVEN TIMM
EXPERIMENTAL COMPUTING FACILITIES
DEPT, FERMILAB**

Outline

2

- I will cover three topics briefly in this talk
 - FermiCloud Infrastructure as a Service facility
 - Scientific Workflows on Cloud in past year
 - Extending site infrastructure to public clouds with the OneFacility project

FermiCloud Resources

3

- **On-demand virtual machine creation**
 - Use cases temporary in nature
 - Try setup, get a fresh virtual machine, repeat.
 - Self-managed, not managed by ECF dept.
- **All FIFE experiments are using this for**
 - Experimental gridftp servers
 - Admin servers for root access on your bluearc and pnfs volumes
- **Some FIFE experiments have used this for**
 - Test monitoring servers
 - Test DAQ servers
 - Virtual machines with unusual memory requirements

FermiCloud Resources

4

- **Documentation**
 - User's guide at
 - <https://cdcv.s.fnal.gov/redmine/projects/fcl/wiki>
 - Request your FermiCloud account if you don't have one via the Service Catalog

Scientific Workflows on Cloud

5

- We have run a significant number of workflows on FermiCloud and Amazon Web Services in the past year
- NOvA far detector cosmic MC
- MicroBooNe high-memory workflow
- Mu2E high-memory workflow
- DES “supercal” workflow.
- Cloud capacity integrated into jobsub and described in the jobsub manual.
- High memory cases will be handled by “partitionable slots” in the main batch system going forward.
- Cloud resources focus on meeting peak demand.

The OneFacility Project

6

- New project that is just starting now:
- Goal to use commercial clouds to meet some of the peak demand.
- Fermilab General Purpose Grid is currently at all-time record usage.
- OneFacility will provide unified interface to local resources, grid sites, commercial clouds, and supercomputer sites.
- Emphasis on getting big data to and from the cloud.

Early OneFacility Use Cases

7

- **NOvA full reconstruction chain (Data and MC).**
 - Estimate 2.1M hours of computing, (20% of what NOvA did last year)
 - 10 MC reconstruction campaigns, 6 data reconstruction campaigns
 - Grant received to fund this campaign
- **CMS Run II MC generation/reconstruction**
 - Estimate ~35M hours of computing, 50000 cores x 1 month.
 - Needed quickly after LHC beam spot is established
- **Build with phased approach**
 - First what we need for these two use cases
 - Then add and modify features based on what we learn.

OneFacility

8

- We need reliable FIFE benchmarks!
 - Is 1 hour of AWS equal to 1 hour on FermiGrid?
 - Early NOvA work said YES but we need to measure it much more accurately than we've done so far.
- We will be coming to FIFE stakeholders looking for more use cases.
- This is a new project, you will hear much more about this at other FIFE meetings.