Fermilab

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

On-demand Computing for Scientific Workflows using Amazon Web Services

Grid and Cloud Services Department, Fermilab

Claudio Pontili

amazon web services | Certified

Solutions Architect - Professional

Supervisors:

Gabriele Garzoglio

Steven Timm

Presentation Outline

- Introduction: why do we want to use Commercial Cloud?
- GlideinWMS: a simple way to access resources
- Running up to 1000 simultaneously jobs on Amazon Web Services and FermiCloud
- Task A: Moving software and data to Commercial Cloud
- Task B: Adding and removing HTCondor to GlideinWMS using AWS
- Task C: Hybrid Cloud solution AWS+Fermicloud
- Conclusions

Introduction: Fermilab Experiment Schedule

- Measurements at all frontiers – Electroweak physics, neutrino oscillations, muon g-2, dark energy, dark matter
- 8 major experiments in 3 frontiers running simultaneously in 2016
- Sharing both beam and computing resources
- Impressive breadth of experiments at FNAL

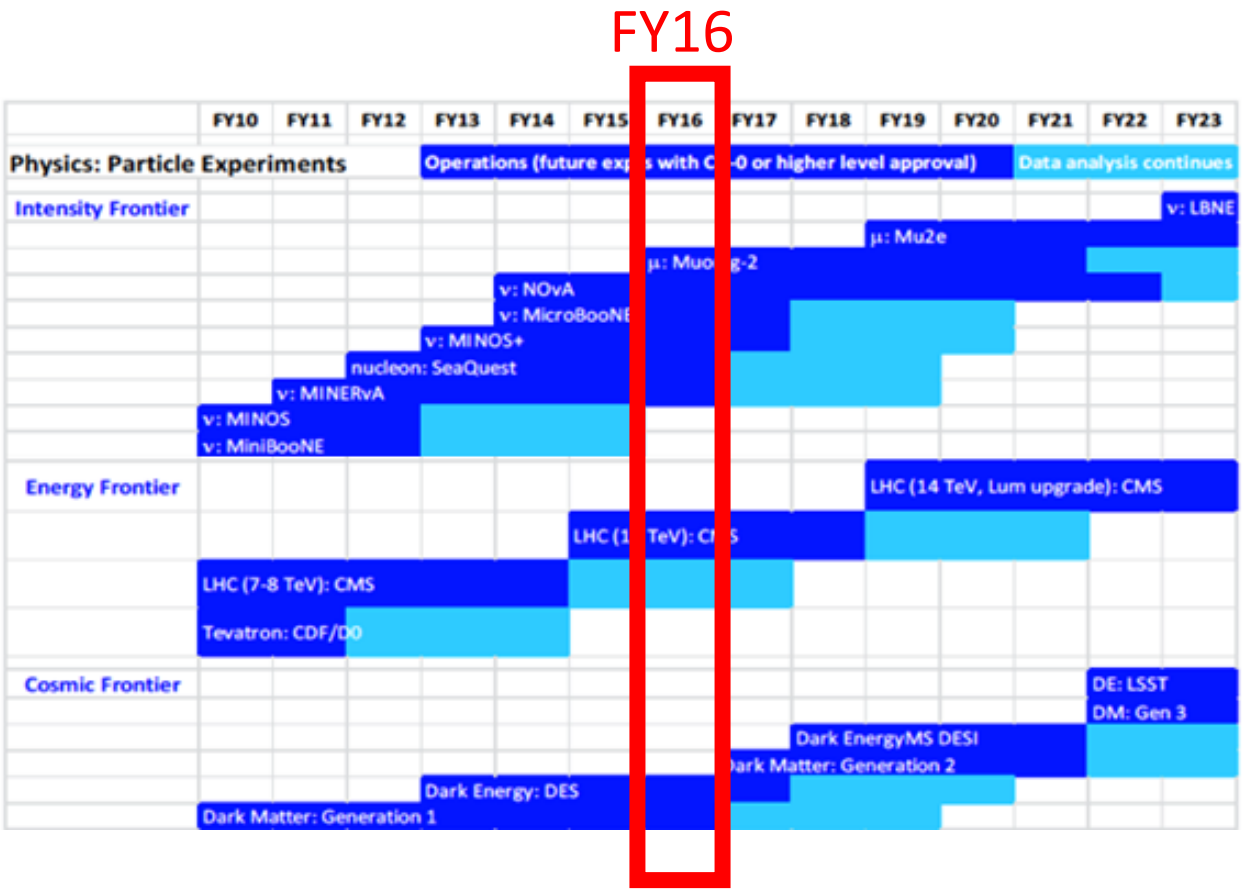

Introduction: Computing requirements for experiments

- higher intensity and higher precision measurements are driving request for more computing resources than previous “small” experiments
- beam simulations to optimize experiments - make every particle count
- detector design studies - cost effectiveness and sensitivity projections
- higher bandwidth DAQ and greater detector granularity
- event generation and detector response simulation
- reconstruction and analysis algorithms

Experiment	Allocation (slots)	Average Utilization (last quarter)		FY15 Requests (slots)
		Average slots	Peak Slots	
ArgoNeut	200	2	164	0
DES	980	64	970	0
LBNE	500	541	5653	1250
Mars	2050	583	4200	0
MicroBoone	500	118	2443	800
Minerva	1600	916	5051	1090
Minos+	1200	2394	8078	2000
Mu2e	1000	1056	4965	3000
Muon g-2	200	129	3638	600
Nova	1300	1363	4969	1800
Total	**9530**	**7166**	**40131**	**10540**

Introduction: Slots Fermilab, OSG, & Clouds

- Current full capacity of FermiGrid → ~30k slots

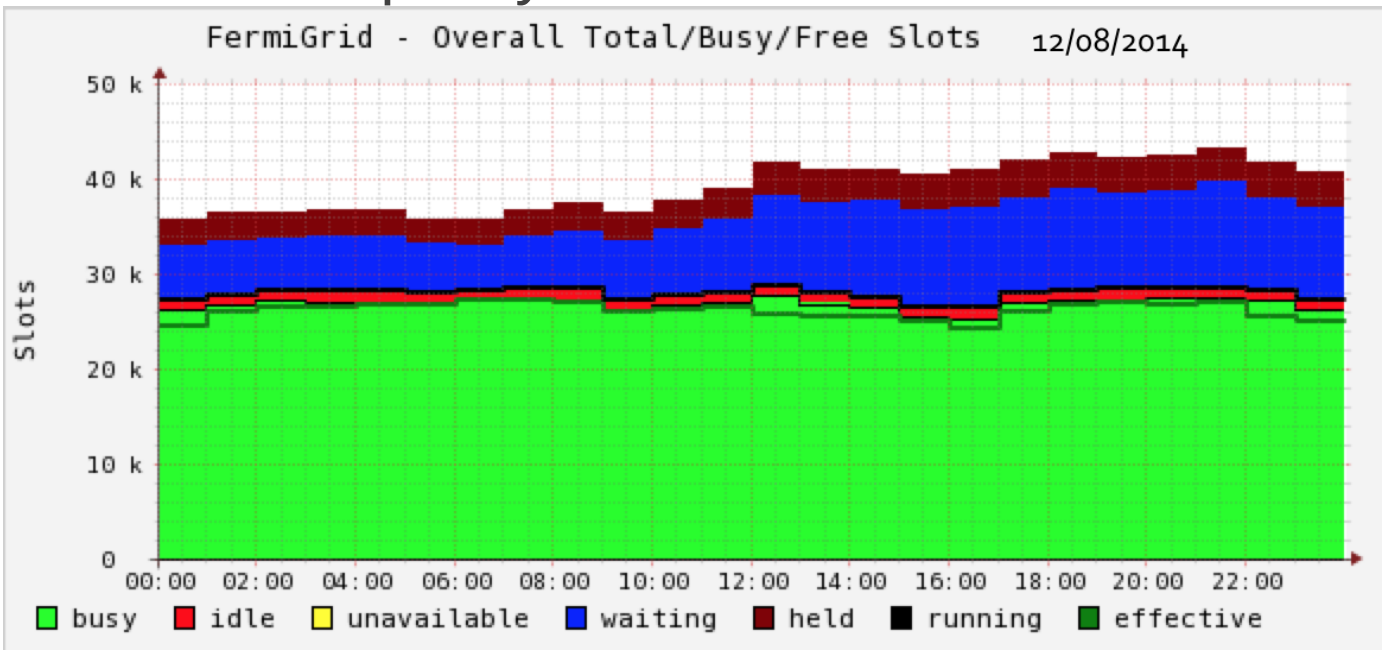

- Full capacity of OSG (Open Science Grid) → ~85k slots
- Additional OSG opportunistic slots → 15k – 30k
- Additional per-pay slots at commercial Clouds

Federation via GlideinWMS – Grid and Cloud Bursting

Unified submit tool for grid and cloud using HTCondor

Users

Jobsub_submit

Job Queue
FrontEnd

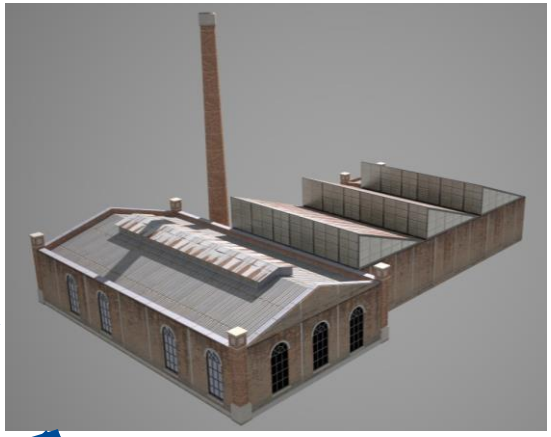

GlideinWMS
Pilot Factory

Jobs

Pilots

FermiGrid

OSG

Amazon
AWS

Gcloud
KISTI

Fermi
Cloud

Fermilab

Running AWS NovA Jobs as function of time, Oct 23. 2014

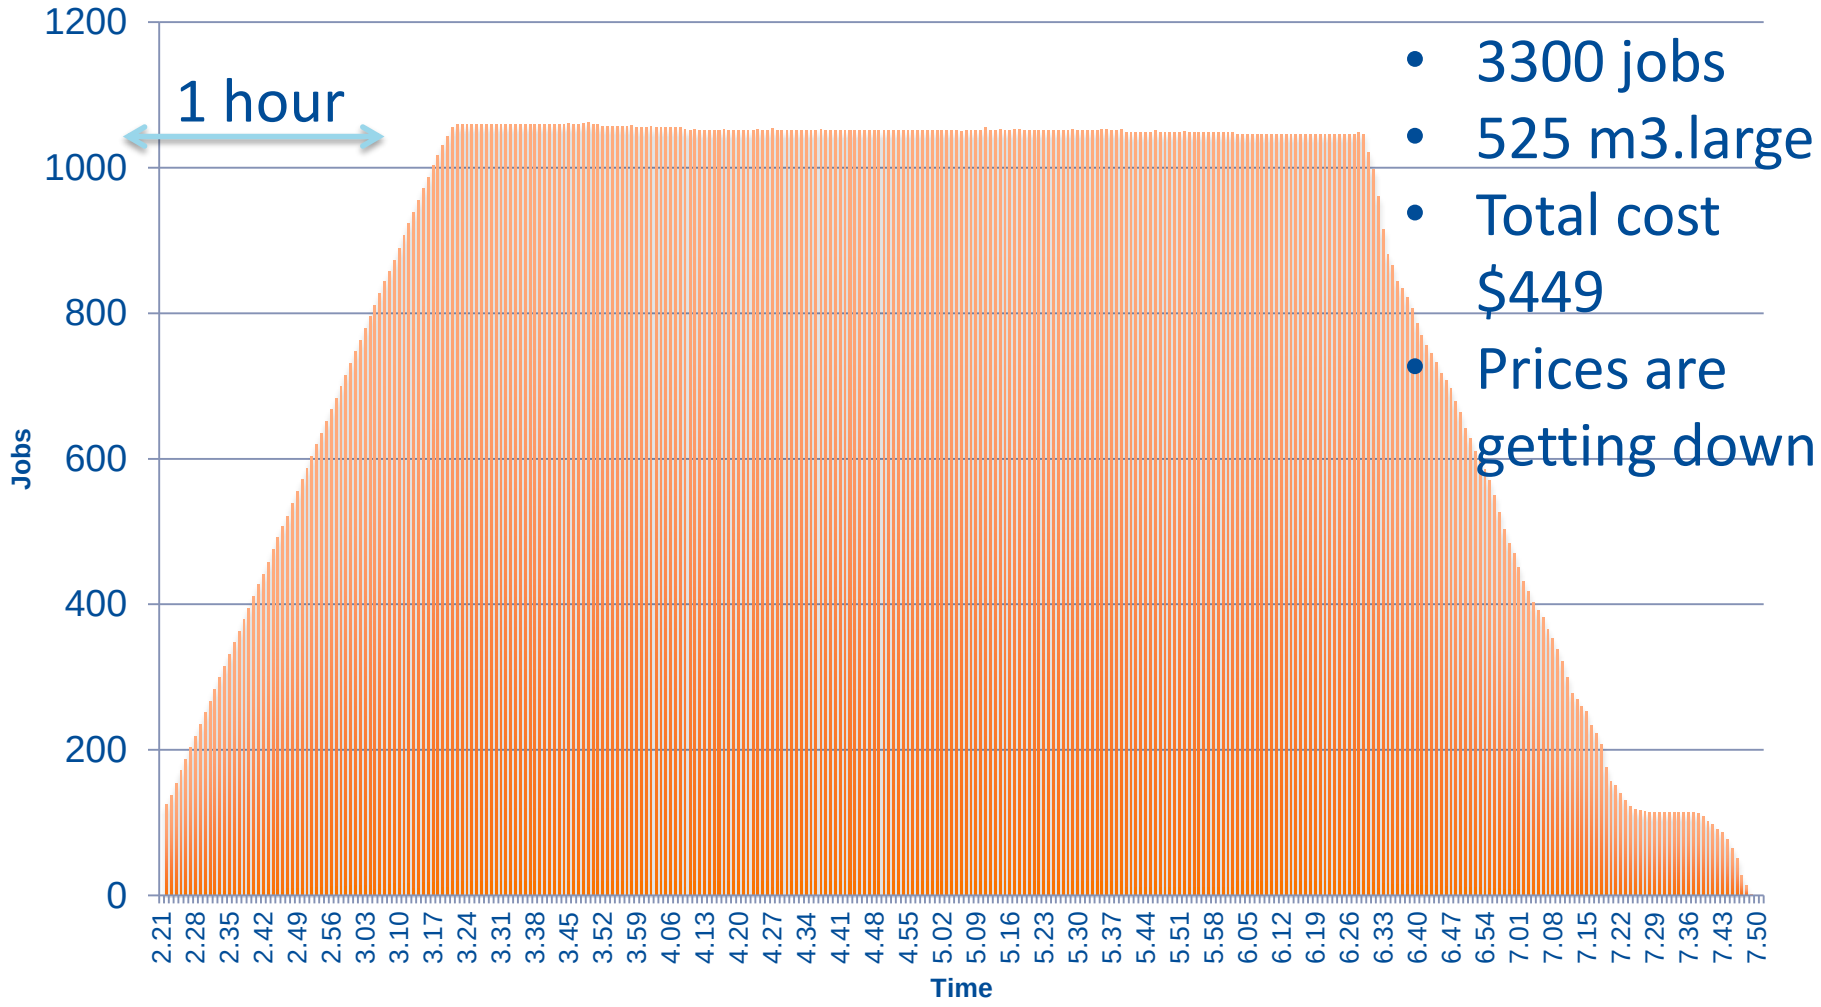

- 3300 jobs
- 525 m3.large
- Total cost $449
- Prices are getting down

Task A: Moving software and data to commercial cloud using Scalable Squid Servers

- Need to transport software and data to the Cloud

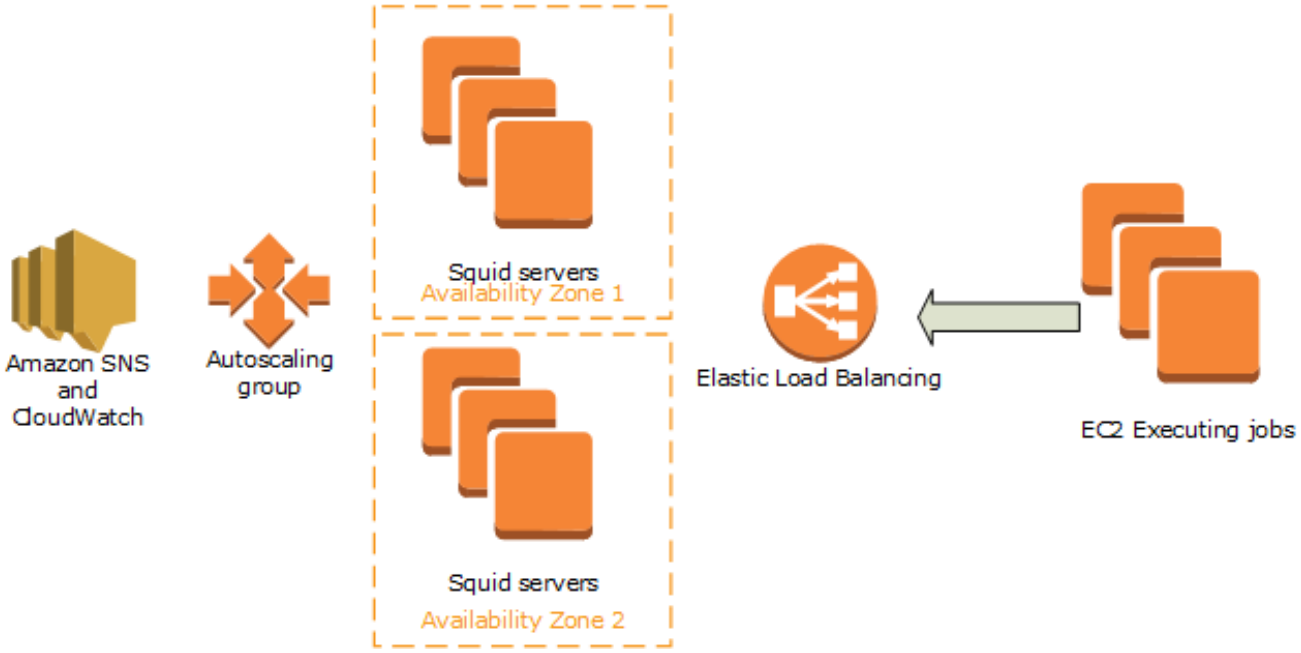

- Auto-scalable squid servers, deploying and destroying in 30 seconds using CloudFormation script

Task B: Auto-Scaling GlideinWMS adding/removing HTCondor using Amazon Web Services

- New resources are made available through the WMS (HTCondor)
- The system is designed to scale by adding servers
- Problem: the submission system is a stateful service
  - Easy to scale up
  - Hard to scale down
- Solution? Manage lifecycle of each server using AWS Hooks and Standby (released July 30th 2014)

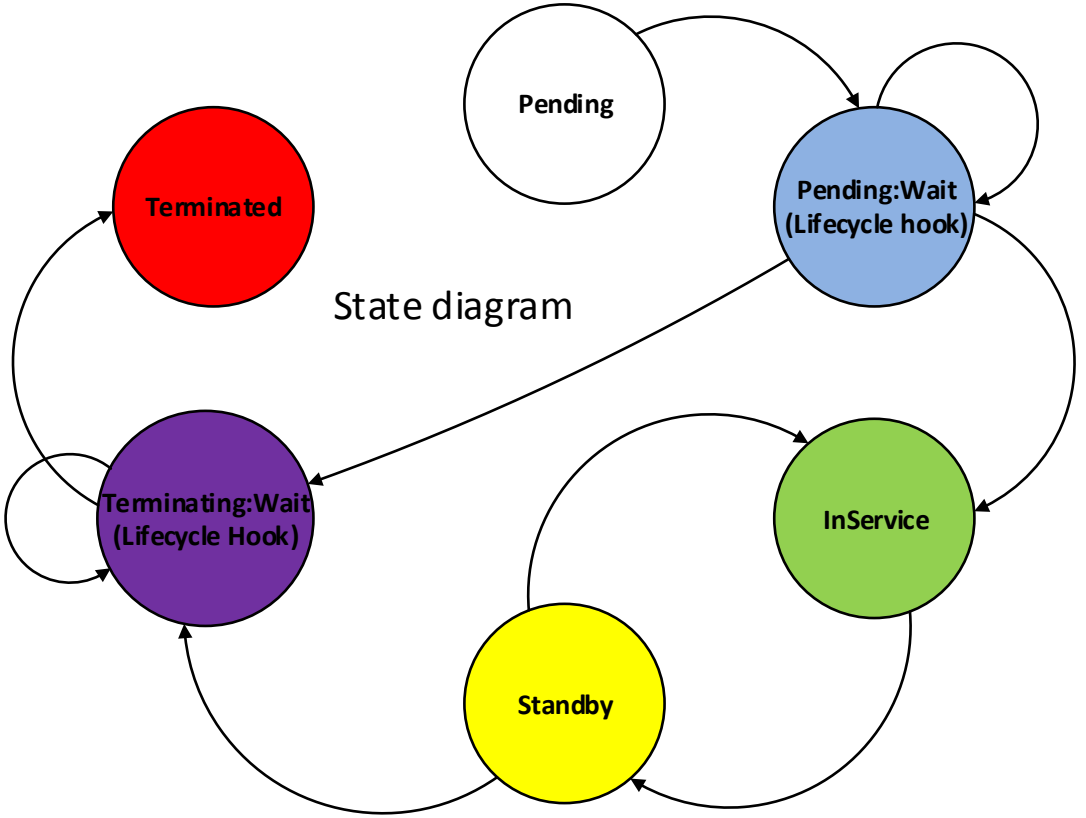

Task B: Auto-Scaling HTCondor using AWS - 2

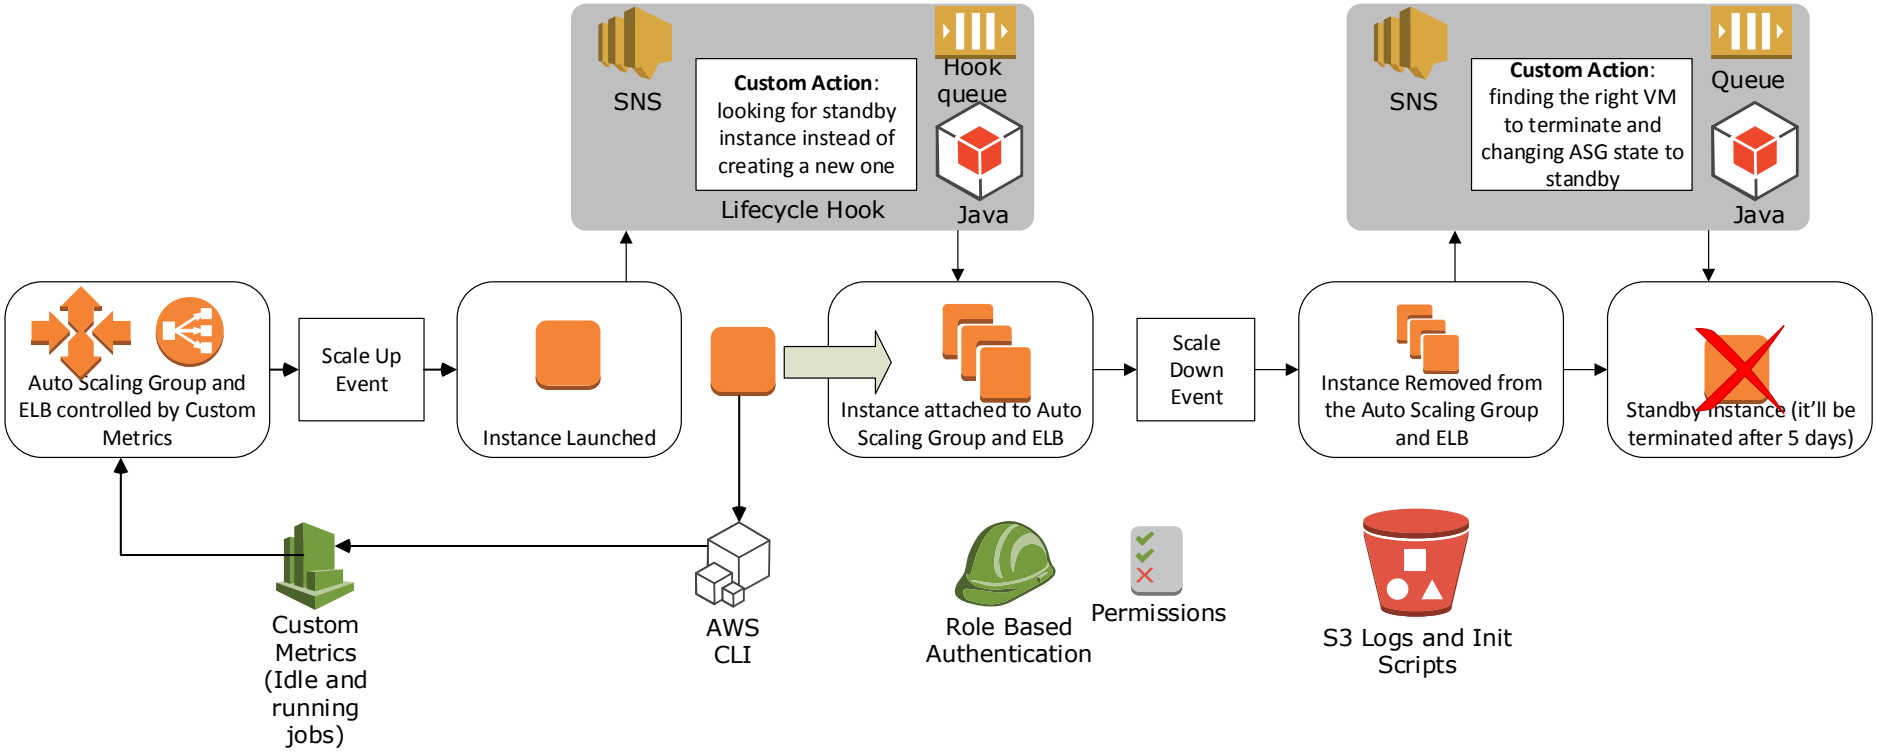

- At least 7 different Amazon services
- 2 different programming languages (Java and bash scripting)
- Role based authentication: auto-generating and auto-rotating logins and passwords

Task C: Hybrid cloud – Fermicloud and Amazon Web Service

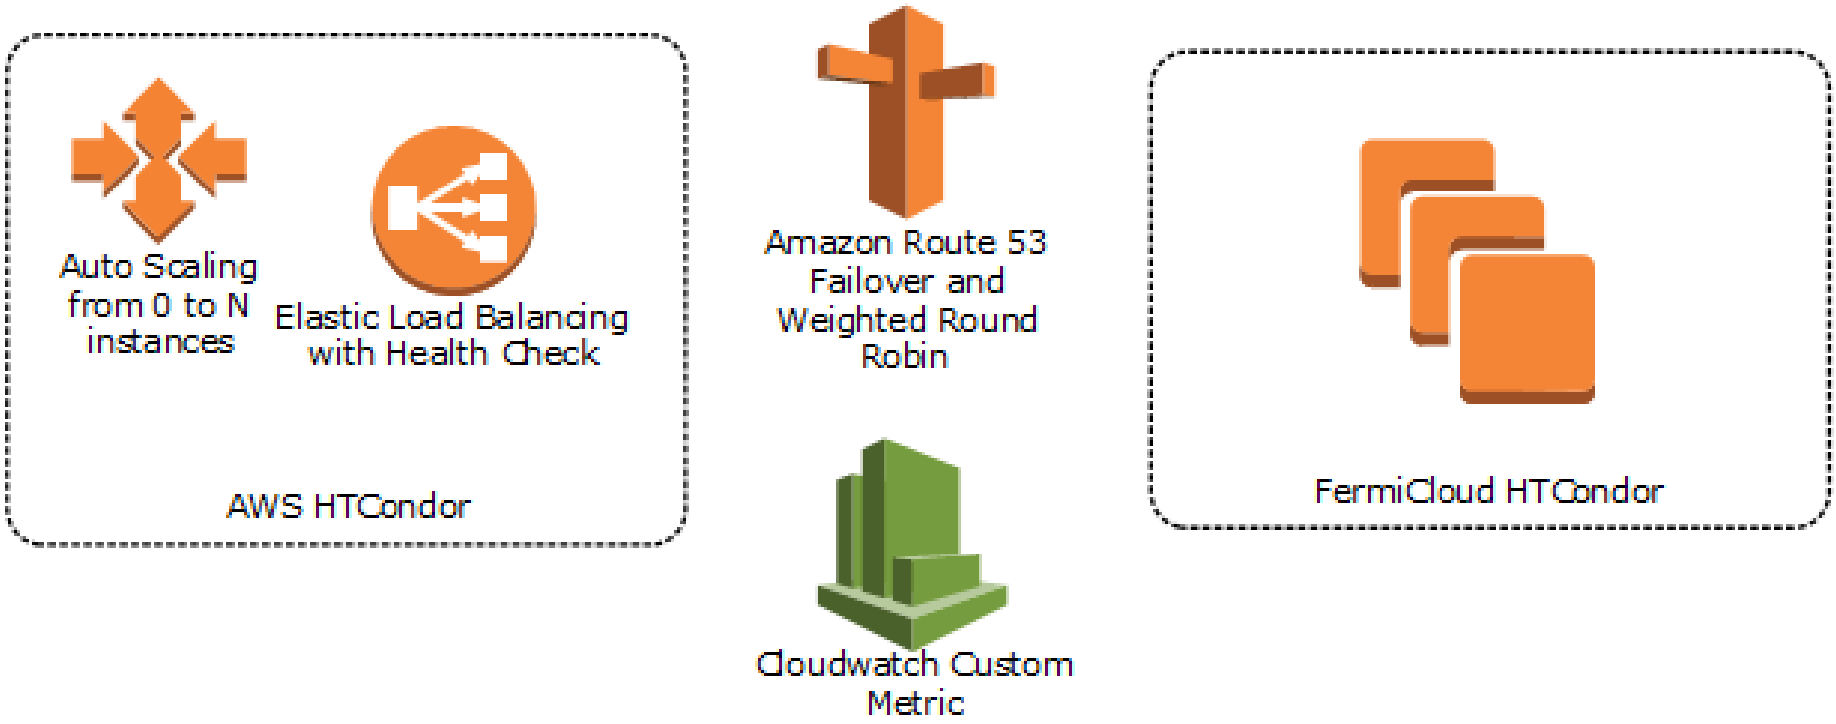

- No change for the final user!!! He just wants to deploy his job in the same way and to computed it as soon as possible
- Now we can handle spikes of traffic using commercial cloud
- We pay AWS only during spikes

Conclusions

- Experiments have an increased need for computing resources with an increased diversity of requirements
- Managing these needs (especially peak demand) is a major focus
- Experiments are being enabled to use a diverse set of resources: Local, Grid, and Cloud
- Need to demonstrate sustainability and cost effectiveness of the commercial Cloud solution for physics use cases (e.g. NOvA MonteCarlo)
- This work demonstrated the scaling of on-demand services in support of scientific workflows using native Amazon Services

Thank you

- Questions?