

Cloud integration with ATLAS Distributed Computing: status, challenges, future work

Fernando Barreiro Megino
WLCG Workshop, Lancaster University
8 November 2022




UNIVERSITY OF
TEXAS
ARLINGTON

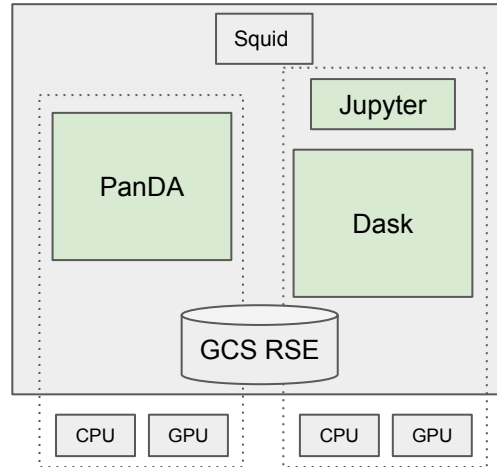


Independent cloud sites

Amazon (CSU Fresno grant)



Google (ATLAS grant)




- Self-contained, cloud-native, vendor-agnostic, auto-scaled, auto-healing
- Rucio:
 - Object Store + Signed URL + http protocol
 - Integrated in Rucio and WLCG MW
- PanDA:
 - Harvester submitting to K8S directly
 - CVMFS K8S plugin
 - Frontier Squid either in K8S cluster or as separate load balanced VM group
 - Auto-scaling: no jobs, no cost

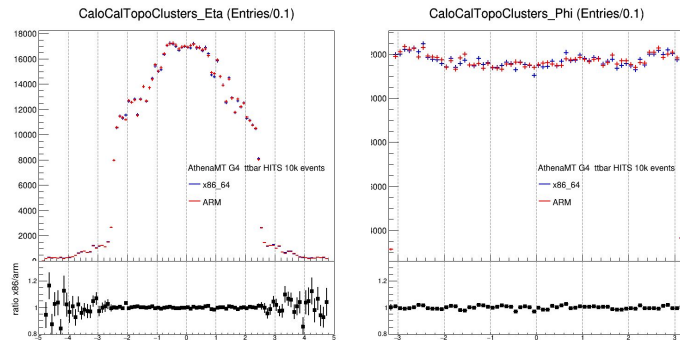
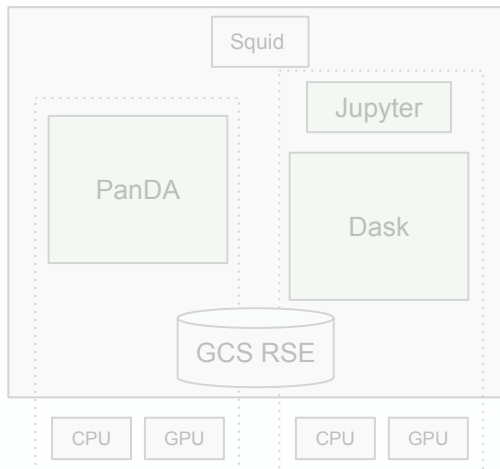
arm64 Physics Validation on Amazon

Preliminary validation
by Johannes Elmsheuser

Amazon (CSU Fresno grant)



Google (ATLAS grant)



Full physics validation [signed off](#) in September

First ATLAS bulk simulation task on **ARM processors** in Amazon




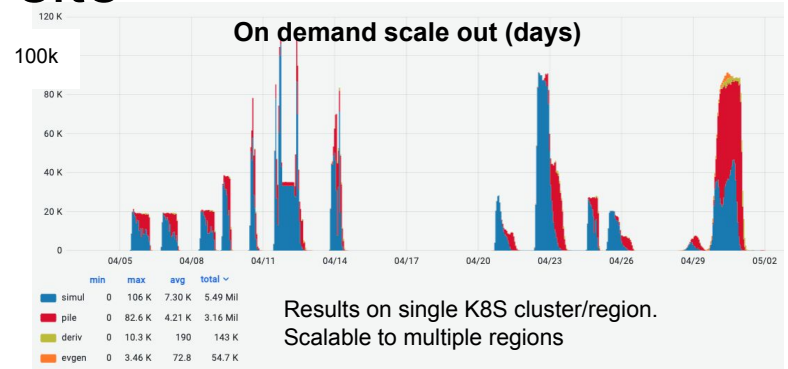
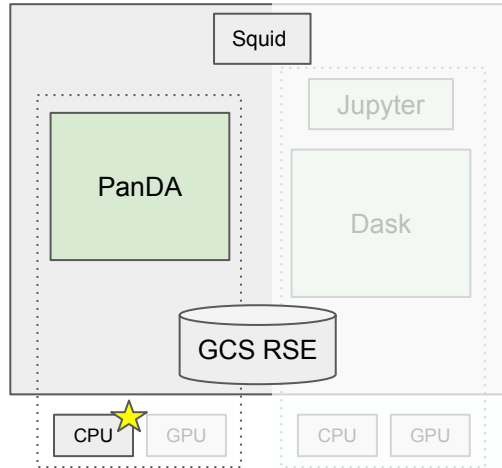
(~4k USD list price)

On demand vs fixed size Google site

Amazon (CSU Fresno grant)



Google (ATLAS grant)

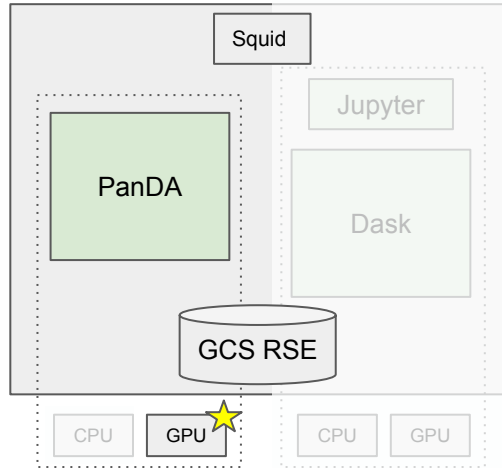


On demand GPU queue on Google

Amazon (CSU Fresno grant)

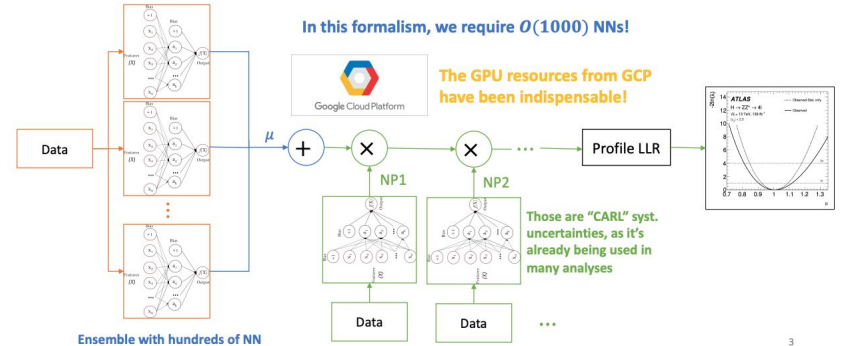


Google (ATLAS grant)




Jay Sandesara:

[Off-shell Higgs Measurement Using a Per-Event Likelihood Method](#)



On demand GPU queue



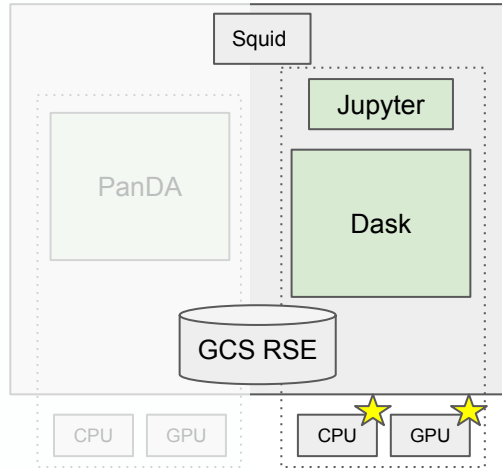
~3k USD list price (not our actual cost)

Jupyter/Dask/TF for user analysis

Amazon (CSU Fresno grant)



Google (ATLAS grant)



Welcome to atlas

Sign in with

Your X.509 certificate

CERN SSO

Not a member?

Apply for an account

You have been successfully authenticated as

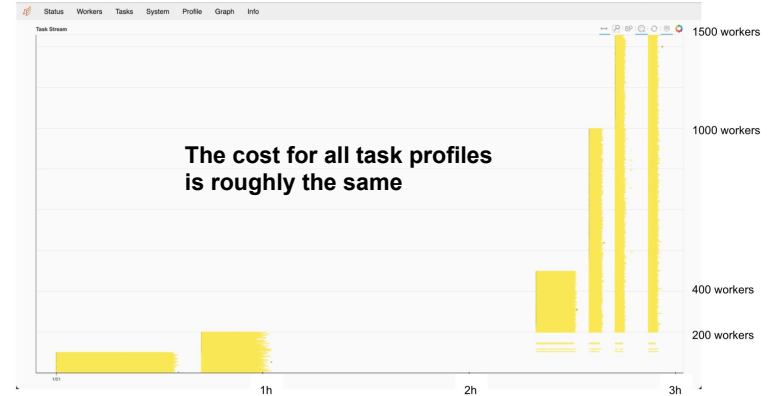
CN=Fernando Harald Barreiro
Megino,CN=643806,CN=fbarreir,OU=Users,OU=Organic
Units,DC=cern,DC=ch

Server Options

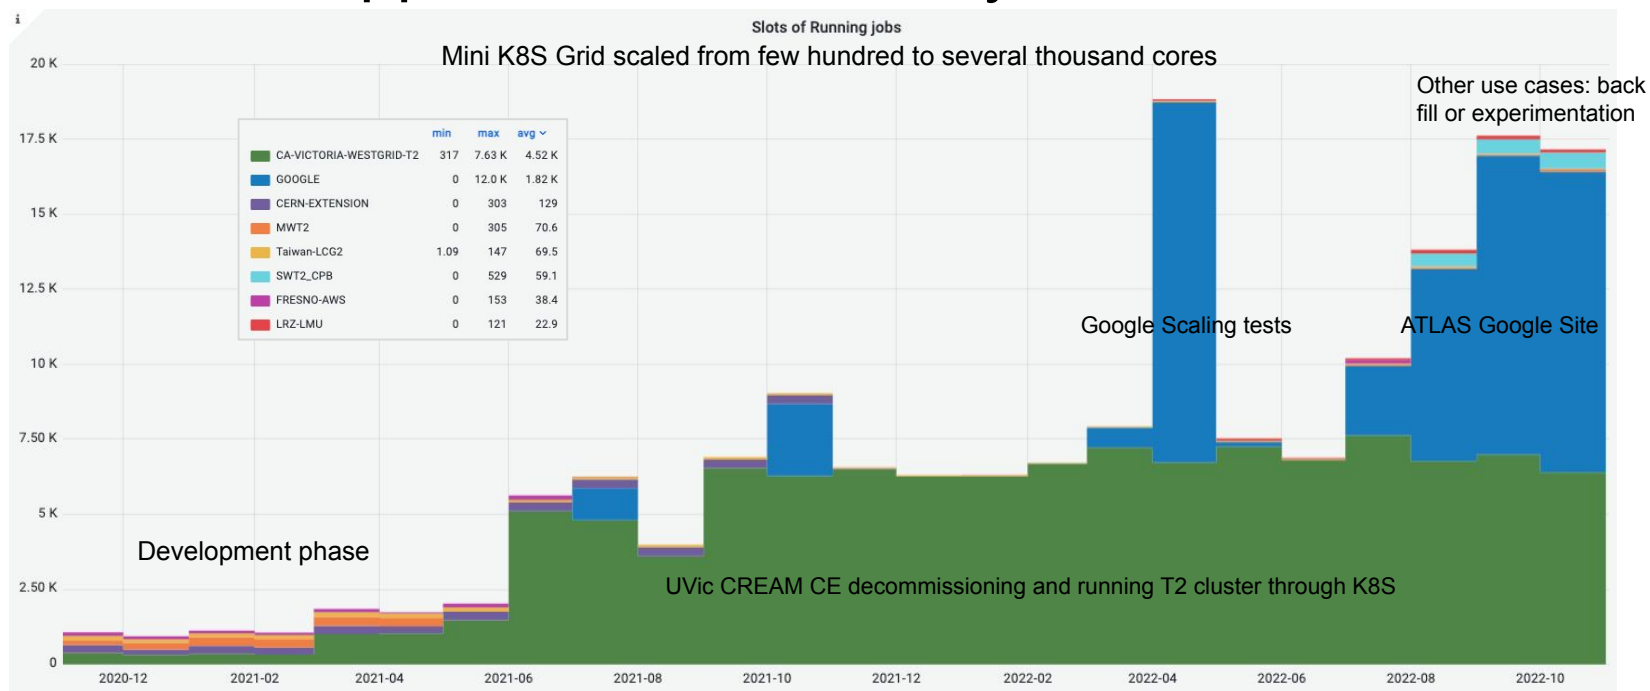
- PHYSLITE environment [prod]
Validated environment for experiments with DAOD_PHYSLITE with uproot, awkward etc. **Columnar analysis**
- ML environment
keras, fastbuffers, joblib, pillow, pytz, scikit-learn, scipy, uproot, root_numpy **ML/AI applications**
- ML environment with T4 GPU
Similar to the previous ML environment, but it boots a VM with a GPU just for YOU. Only to be used when working on ML-GPU applications. Setting up your notebook will take 10 minutes or more. **+GPU**
- ML environment with P100 GPU
Similar to the previous ML environment, but it boots a VM with a GPU just for YOU. Only to be used when working on ML-GPU applications. Setting up your notebook will take 10 minutes or more. **+GPU**
- ML environment with K80 GPU
Similar to the previous ML environment, but it boots a VM with a GPU just for YOU. Only to be used when working on ML-GPU applications. Setting up your notebook will take 10 minutes or more. **+GPU**

Start

Dask scale out exercises: effect of cluster size on task duration. Resources are booted up as necessary
(By Lukas Heinrich)



PanDA-K8S approach works on any cluster



University of Victoria



arbutus cloud



Google Cloud



CERN



THE UNIVERSITY OF CHICAGO



UNIVERSITY OF TEXAS ARLINGTON



FRESNO STATE



aws



lrz Leibniz-Rechenzentrum der Bayerischen Akademie der Wissenschaften

Challenges

- Grid compatible integration, e.g. using nested containers
- CVMFS K8S plugin: needed experience until it got stable
- What do we miss out by not running a CE?
 - APEL Accounting
 - Ignored for commercial clouds
 - [KAPEL](#): developed by Ryan Taylor for UVic reporting
 - Authorization
 - Manual generation and transmission of Kubeconfig files into Harvester

Challenges (2)

- IGTF CA Certificates
 - IGTF=Interoperable Global Trust Federation
 - Google/Amazon CAs do not belong to IGTF
 - Grid sites do not trust them and 3rd party transfers do not work
 - Workaround to enable 3rd party transfers
- Direct IO on Object Store
 - Object stores don't support [multirange requests](#). Potential workaround to get direct IO working
- LHCONE integration
 - So far data travelling through cloud network and internet has worked fine
 - Is it possible to add a commercial cloud (very broad IP range)?
 - How to guarantee only LHC data is travelling?

Future work

- Demonstration of elasticity: combine flat and dynamic resources in a controlled way
- Access to “exotic” resources and gain experience
- Egress cost optimization and potential evaluation of better networking models
- Total Cost of Ownership evaluation
- Possibility for users to get direct access to cloud and explore particular services (e.g. AutoML)


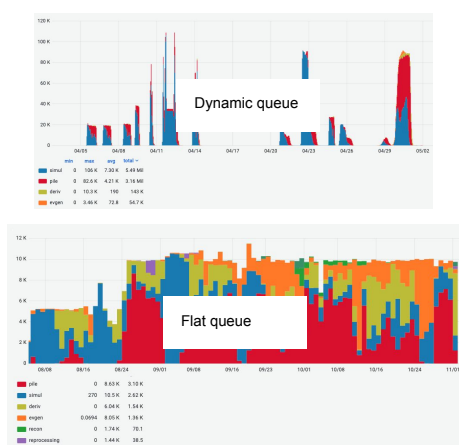


Image credits: NVIDIA GTC May 2020 Keynote

Questions