ATLAS/Google R&D

Year 2 Program Planning

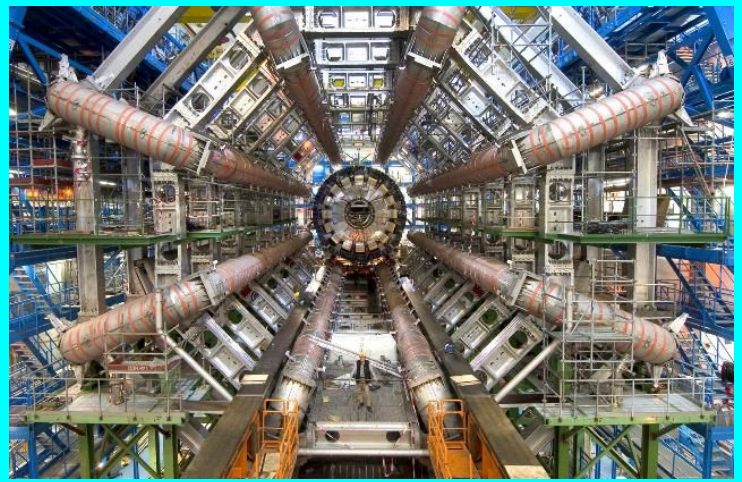

February 17, 2021

Agenda

- Program Challenges and FY 2020 Accomplishments
- US ATLAS R&D Year 2 Plans - Milles, Kaushik/Alexei
- Physics Analysis Use Cases and GCP - Lukas/Nikolai
- GCP Infrastructure for R&D - Initial Considerations & Questions - Fernando
- Data Management - Cedric
- Next Steps

HL-LHC Challenges

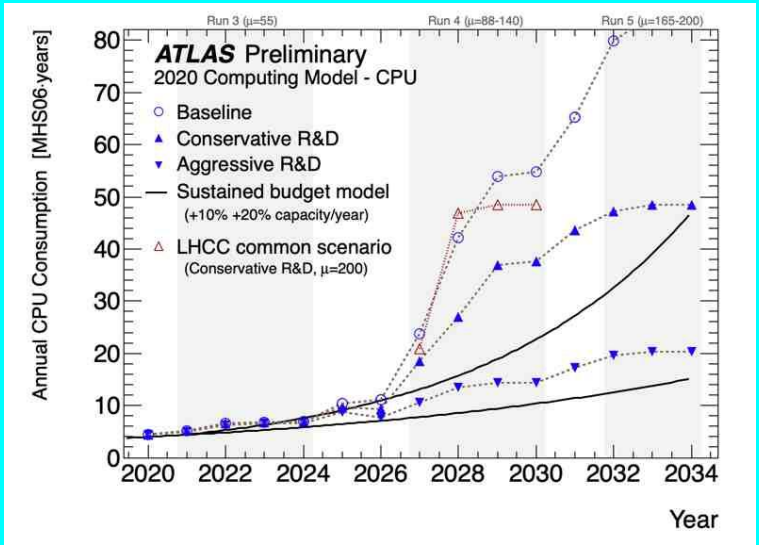

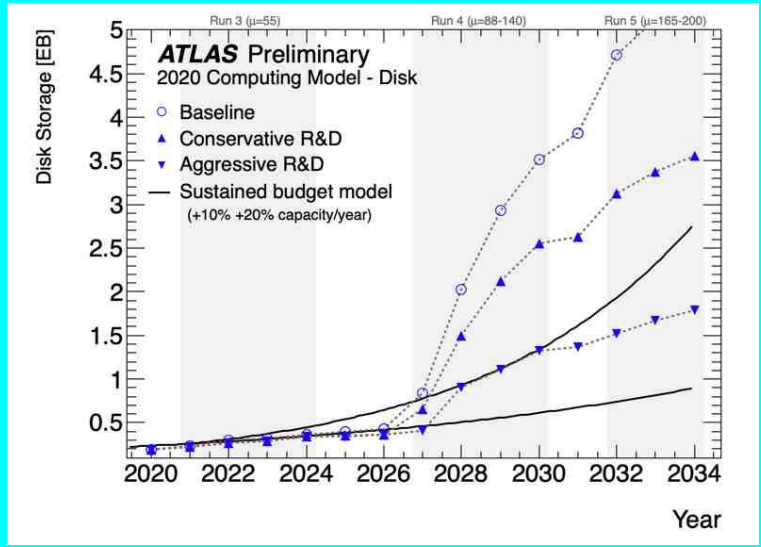

- Future computing needs for HL-LHC are well above flat budgets
- CPU usage, storage needed, and networking - all will require optimization and innovation in future computing models in preparation for Run 4

FY 2020 R&D Key Accomplishments

November 1, 2019 - October 30,2020

- Proved that Google Cloud Storage (GCS) & GCP can integrate with the PanDA + Rucio workflow
- Proved that GCS is a reliable and easy to access alternative to the present mass storage (disk + tape on the grid)
- Created formal structure for US ATLAS and Google personnel to effectively work together to address technical challenges and identify ways to leverage cloud-native tools in GCP
- Identified issue with cloud egress costs for large scale data downloads from GCP. Used this information to address this gap in recently awarded DOE Google Enterprise Agreement, which now offers a free egress model
- US ATLAS and Google team worked together to develop object-storage friendly data formats for Root data
- Began exploring methodologies to execute data analysis within GCP, leveraging native GCP analytics tools
- Worked on ML operations specifically for trace reconstruction

Collaborators*:

- ATLAS CERN : ~0.25 FTE
- US ATLAS : ~1.25 FTE
- Google : ~1 FTE

Tracks (from 2019 DOE white paper)

- Track 1: Data Management Across Hot/Cold Storage
- Track 2: Machine Learning, TPU vs. GPU for GNN training
- Track 3: Optimized I/O and data formats for object storage
- Track 4: End user analysis conducted worldwide at PB scale

Year 2 R&D Overview

- **Plans**
  - Continue US ATLAS / Google partnership designed to explore forward-looking **architecture** that can **scale** with **growing storage & compute demands**
  - Identify potential efficiencies delivered by cloud computing allowing users to spend **less time on infrastructure** management & **more time on science**

- **Objectives**
  - Work with US ATLAS researchers to explore benefits of leveraging research tools (AI; ML; data analytics) in GCP
  - Test end-to-end analysis on GCP with frequently used datasets to identify additional benefits to scientific community
  - Conduct large scale prototyping to ensure GCP can scale to meet US ATLAS compute and storage requirements. Increase data upload to GCP and prove methodology for doing analysis in the cloud, versus at the local level
  - Address remaining challenges with integration of US ATLAS management tools (PanDA+Rucio) w/ GCP
  - Google team to provide knowledge transfer to US ATLAS users with focus on becoming more proficient with GCP

- **Resources**
  - GCP Credits/Compute Cycles
  - Access to Google SMEs
  - Google to offer training and technical talks to upskill users on GCP tools and capabilities

GCP Credits

- Year 2 ATLAS prototyping and other work on GCP will be supported by GCP credits
- Multiple Sources of GCP Credits:
    - ATLAS funded = $250K+/-
    - Google Funded = $80K+/-
    - DOE HQ Funded = $30K+/-
- GCP credits will allow users to access any and all services within GCP
- Program team will work with users to set up “guard rails” to ensure the available number of credits is not exceeded
- Team should discuss any projects/tests that might use significant GCP resources so that group is in alignment regarding how GCP credits will be used
- As part of the Google Subscription Agreement (SA) model, the team will have some flexibility to exceed the total amount of credits allotted; however, this will need to be managed tightly with Google support

Next Steps