Dark Energy Survey Experience with Gravitational Wave Follow-up Workflows on HEP Cloud

G Garzoglio[1], K Herner[1,†], S Timm[1], and A Tiradani[1]

[1] Fermilab

[†] Corresponding author: kherner@fnal.gov

1 Overview

The Dark Energy Survey Gravitational Wave (DES-GW) program aims to use DECam, located on the 4-meters Blanco telescope at the Cerro Tololo Inter-American Observatory in Chile, to perform rapid optical follow-up of gravitational wave (GW) triggers from the LIGO-Virgo Collaboration. Once a gravitational wave trigger is received, the DES-GW team will calculate an observing plan to search for electromagnetic counterparts to GW events. Such counterparts can appear in GW events caused by several methods, including the merger of two neutron stars, or a neutron star and a black hole.

The DES-GW team searches for electromagnetic (EM) counterparts by using the Difference Imaging pipeline (Diffimg), which takes previous images of the same region of sky as the search image ("template" images) and subtracts the ensemble from the search image. Then the pipeline scans for objects left over after the subtraction; these "new" objects are further analyzed to see if they are consistent with models of EM counterparts to GW events. If any candidates are flagged as possible counterparts, the DES-GW can send a communique to other telescopes involved in the LIGO-Virgo GW follow up program, so that they can make more detailed observations.

2 Report

The light from electromagnetic counterparts fades on a timescale of days, so it is critical to process each night's images as quickly as possible to determine if there are any counterpart candidates that should be flagged for further observation. The DES-GW goal is to perform a complete processing of all images from a given night within 24 hours. A typical observing night will produce between 100 and 200 images. Each search image requires a preprocessing step that takes 2-6 hours and then 60 parallel Diffimg jobs, one per CCD, each lasting about one hour. The Diffimg jobs consume a single CPU and use up to 2 GB of memory. After the Diffimg jobs are complete, some additional post-processing takes between 1 and 2 hours per image. One difference in the AWS workflow relative to the standard workflow is that the post-processing would also occur within AWS instead of on an interactive machine, in order to reduce the data egress costs from S3. Therefore, the total resource requirements of following up a GW trigger on a given evening are between 10,000 and 15,000 CPU-hours over a time period of several hours. New images arrive approximately every three to four minutes.

For the HEPCloud test, we begin with a single Diffimg pipeline job for one CCD of an image taken as part of the follow up campaign to the first GW event, GW150914. The Diffimg pipeline code itself was unchanged for the AWS tests, but we set it up so that all output data went to Amazon S3 instead of the Fermilab dCache. The single test job was successful, with outputs being copied to the correct S3 bucket. In the next phase of the test we sent 60 parallel Diffimg jobs, one for each CCD in a single exposure. All jobs completed successfully and stored approximately 1.5 GB of valid output to S3. The total amount of data stored to S3 was approximately 100 GB, with less than 4 GB of the total output needed for the post-processing analysis. The cost of the overall demonstration was $178, with the main costs being $154 for computing services (EC2), $6 for storage (S3), and $17 for AWS support.

These tests prove that the DESGW workflow can successfully run on AWS.