

Pitt-Google write up from the Low-latency alerts & Data analysis for Multi-messenger Astrophysics workshop

Pitt-Google broker will be a full-stream alert broker for LSST. We are currently running the ZTF alert stream and planning to run gravitational wave alert streams. We add value to the alerts (e.g., cross matches) and publish/store to the Google Cloud Platform services Pub/Sub (streaming message service), BigQuery (data warehouse with SQL access) and Cloud Storage (file store). All data is available to pull/query in real-time via APIs, with latencies between one and a few seconds. During the MMA workshop, I (Troy Raen) presented a process by which users will be able to listen to a gravitational wave stream and trigger their own real-time analysis of LSST data in search of an event's electromagnetic counterpart. The search can look both forward and backward in time. Users can compute next to the data with minimal infrastructure management using (e.g.,) Google Cloud Run which includes automatic scaling and fault tolerance.

Researchers are often concerned about the cost of cloud computing, and similar concerns were raised during this workshop. I presented a cost estimate for my example user-analysis process which totaled approximately \$800 per year. This allows for a relatively large amount of data every night: processing 10^5 LSST-sized alerts, querying 250 GB of data, and egress of 10^3 LSST-sized alerts (US to Europe). Cloud computing can greatly reduce the workload required to implement scientific analyses since the majority of infrastructure management can be handled by the cloud provider. However, organizational budgets will need to evolve. Currently, many are scoped to pay their own employees to perform these tasks, but are not scoped to pay a cloud provider.