



Building a 100% cloud-based alert broker: Pitt-Google

Troy Raen
University of Pittsburgh and the Pitt-Google Collaboration

LineA Webinar | December 2, 2021

Pitt-Google Alert Broker Collaboration



Troy Raen

Lead Developer
Grad Student
University of Pittsburgh

on behalf of...



Michael Wood-Vasey

PI, Associate Professor
University of Pittsburgh



Ross Thomson

Solutions Architect
Google



Christine Mazzola Daher

Grad Student
University of Pittsburgh



Daniel Perrefort

Research Assistant Professor
University of Pittsburgh

Outline



1. **Pitt-Google broker's motivation**
2. **Comparison of broker approaches and basics of the Cloud model**
3. **Pitt-Google broker architecture and the Google Cloud services we employ**
4. **Other Google Cloud tools of interest**

Learning Objectives



1. Understand the Cloud model and cost/benefit tradeoffs.
2. Describe two Google Cloud services, and
 - a. how Pitt-Google is using them to develop an LSST alert broker; and/or
 - b. how you might use them with the LIneA IDAC

Pitt-Google broker's motivation

Pitt-Google Broker Motivation



Get the data into Google Cloud,
so that anyone and everyone
can access however much they want,
and do whatever they want with it.

Pitt-Google Broker Motivation



Get the data into Google Cloud,
so that anyone and everyone
can access however much they want,
and do whatever they want with it.

The LSST project is not scoped for this.

Pitt-Google Broker Motivation



**Get the data into Google Cloud,
so that anyone and everyone
can access however much they want,
and do whatever they want with it.**

Caveats:

- Internet access required (but no other hardware, software, etc.)
- Pay for what you use

Pitt-Google Broker Motivation



Get the data into Google Cloud,
so that anyone and everyone
can access however much they want,
and do whatever they want with it.

Goals:


1. Enable scientific collaboration.
2. Expose the data publicly at each step of the processing pipeline, so users can access data that is pre-processed to a level of their choosing.
3. Use Google Cloud *services* to maximize efficiency and minimize the effort required to develop, run, and support an alert broker at LSST-scale.

Comparison of broker approaches and Basics of the Cloud model

Other brokers also run “in the cloud”...

Many other astronomical community alert brokers are also operating in a cloud environment.

This may mean anything from



Other brokers also run “in the cloud”...


Many other astronomical community alert brokers are also operating in a cloud environment.

This may mean anything from

Fink, operating on VirtualData OpenStack-based cloud, cites **benefits of cloud computing** as:

“scalability of processing and cost effectiveness, fault tolerance, shareability of results, as well as increased portability and reproducibility”

(Möller et al., 2020)



Other brokers also run “in the cloud”...



what makes Pitt-Google different?

A fundamentally different broker model based on

Google Cloud *services*

and

***event-driven* processing.**

Cloud Services



In cloud computing, users rely on the cloud provider to manage physical hardware and support access.

The level of support can vary greatly. Google Cloud examples:

Compute Engine (VMs):

- User writes the application code and manages the VM environment, containerization & deployment, scaling to meet demand, fault tolerance, etc.
- User has a lot of control, and a lot of responsibility.

Cloud Functions (serverless compute):

- User writes a function and deploys to the Cloud. Google manages everything else.
- User trades control and responsibility for ease-of-use.

They are both *services*,
but infrastructure management responsibilities are
split between Google and the user very differently .

Cloud Services



The Pitt-Google broker takes advantage of managed services, not VMs, as much as possible.

This results in a lightweight, highly modular, nearly serverless pipeline.

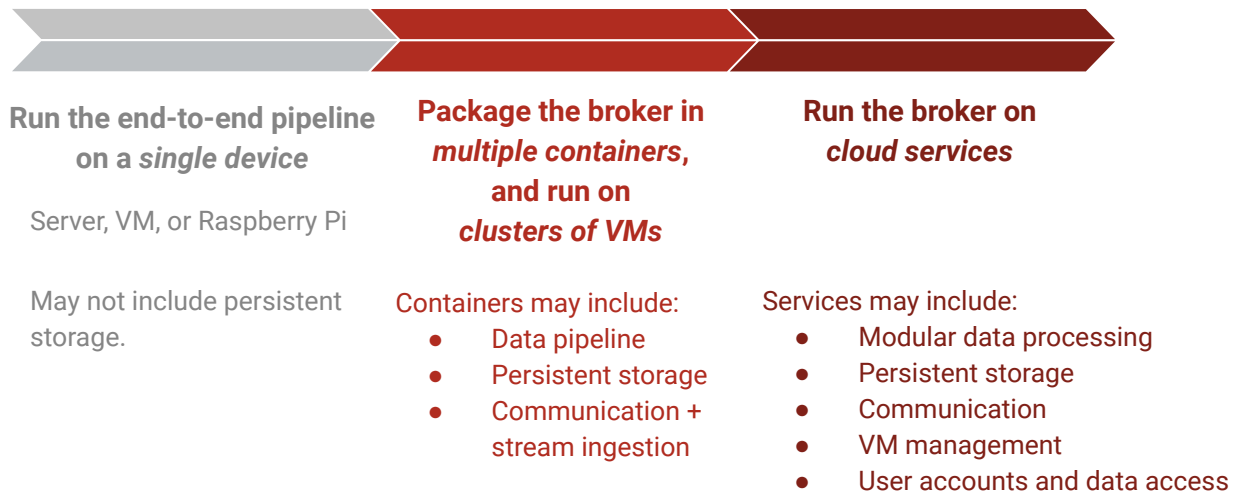
and no direct data storage or delivery responsibilities.

Jargon: Serverless

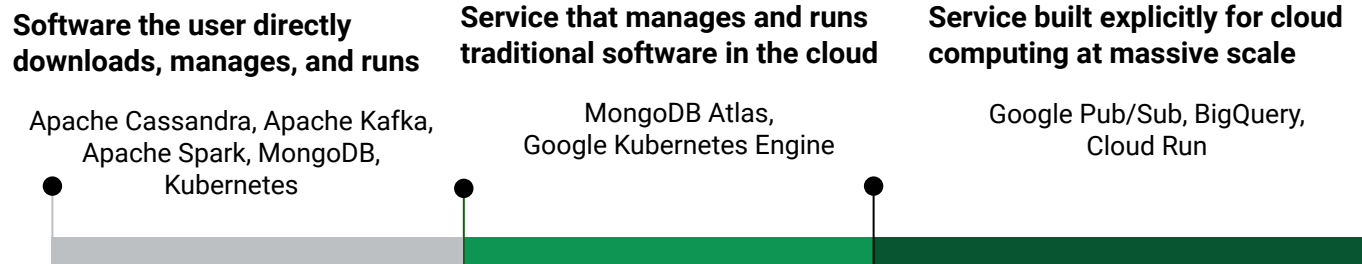
“*developers* of serverless applications are *not concerned with capacity planning, configuration, management, maintenance, fault tolerance, or scaling of containers, VMs, or physical servers*”

https://en.wikipedia.org/wiki/Serverless_computing

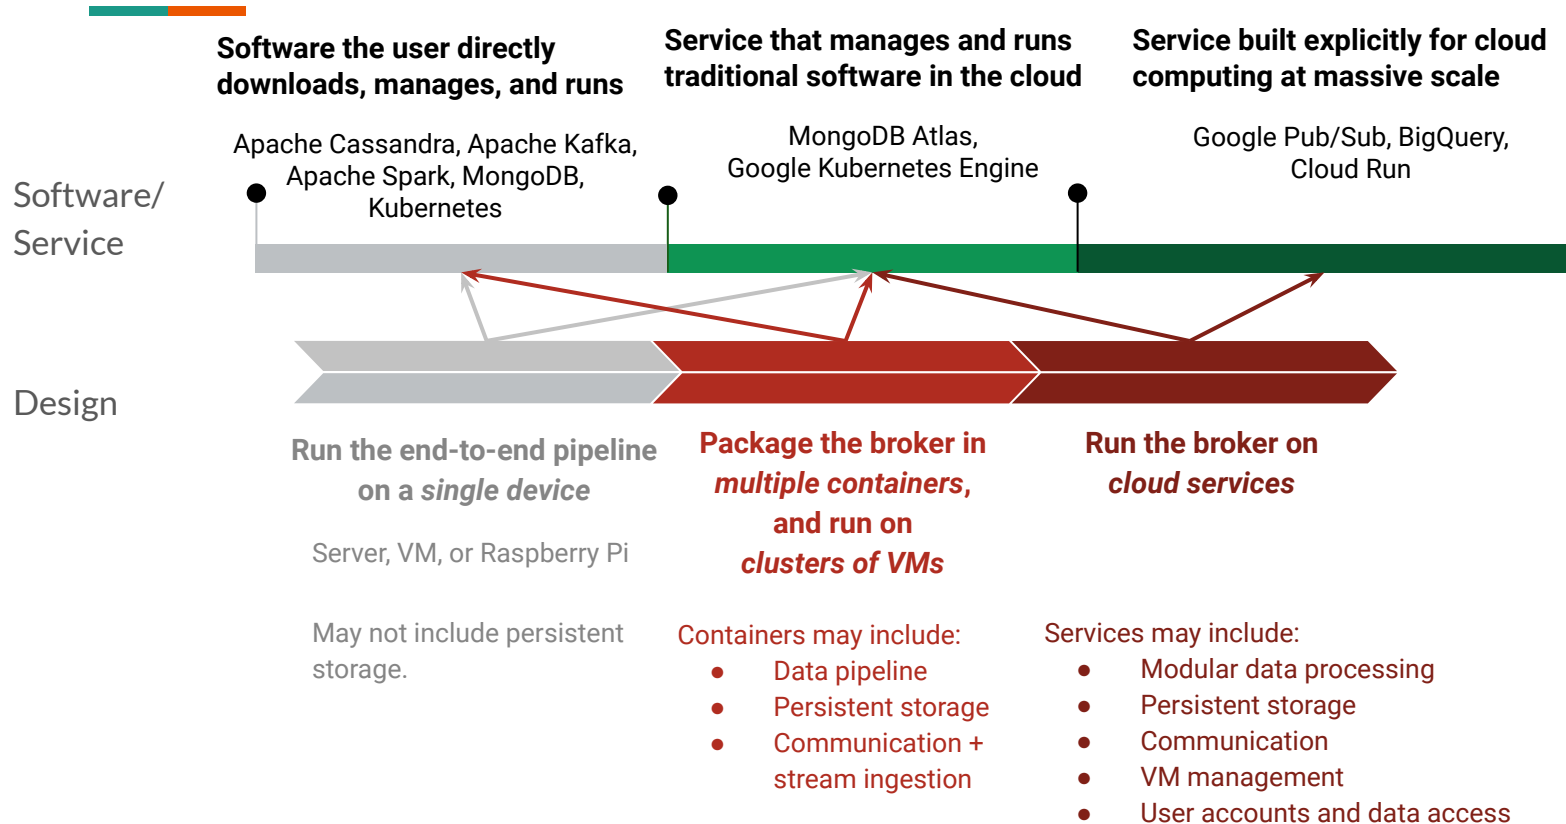
Broker high-level design models



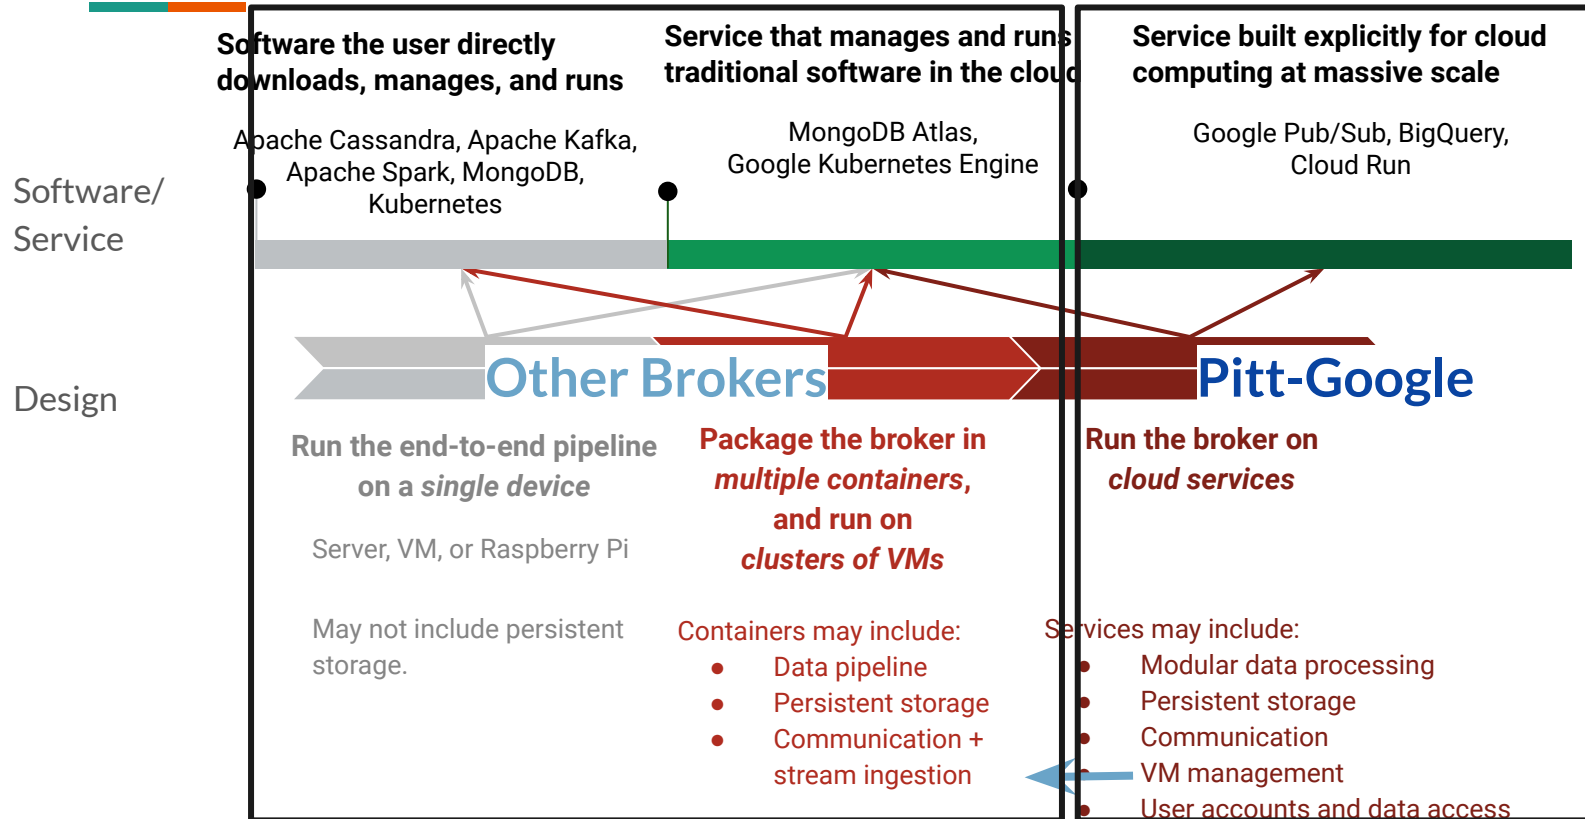
Types of software and services brokers use



Relationship between design choice and software/service



Relationship between design choice and software/service



Data Access from Google Cloud Services

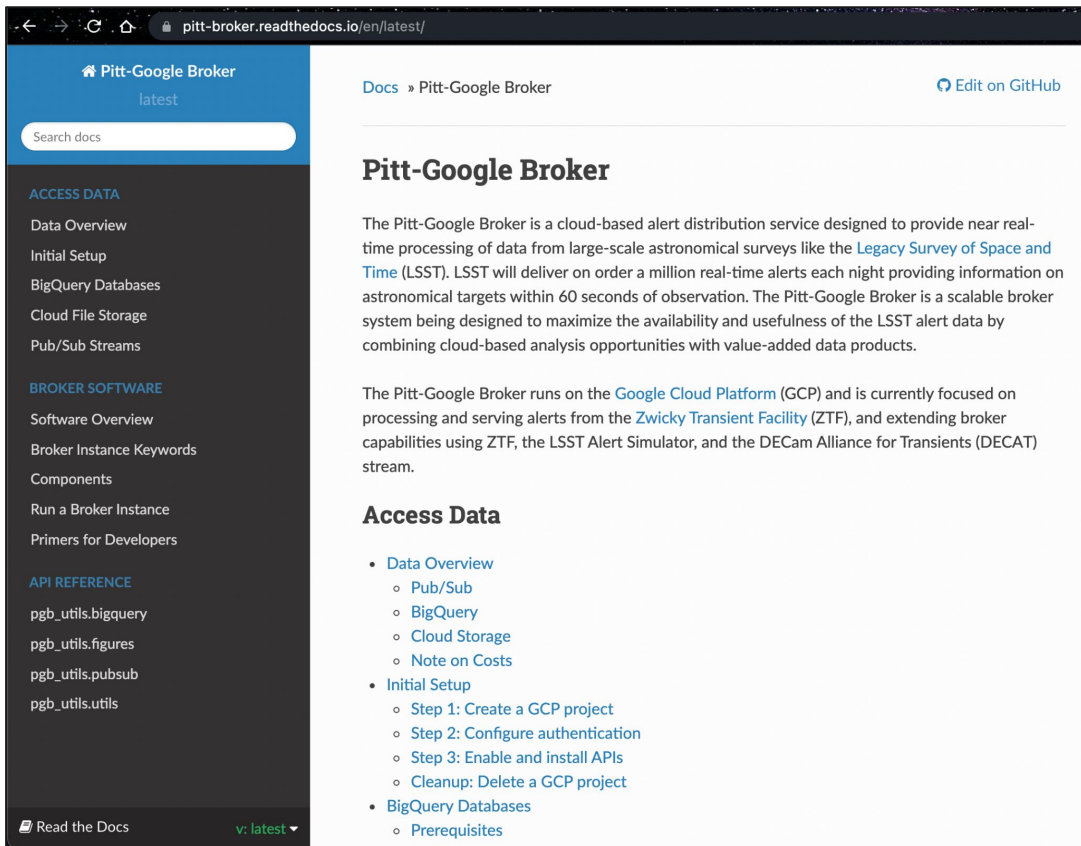


- Pitt-Google stores data in Google Cloud services (BigQuery, Pub/Sub, and Cloud Storage).
- Google manages the data storage, data delivery, and user accounts & credentials.
- To access data, one first creates an account and obtains credentials from Google Cloud.
- By contrast, other brokers manage their own data storage/delivery and user accounts/access. Their users contact the broker for data access, even if the broker is running in the cloud.

How? Methods to access data from Google Cloud:

1. [APIs](#) developed and supported by Google. Access from anywhere, in or out of Google Cloud.
2. Within Google Cloud, many of their tools and services are integrated, making it easy to write streaming data pipelines or compute/analyze data at the source, without having to think about servers or data management.

Data Access from Google Cloud Services



The screenshot shows the documentation page for the Pitt-Google Broker on Read the Docs. The page title is "Pitt-Google Broker" and it includes a search bar, a navigation sidebar, and a main content area. The sidebar lists sections like "ACCESS DATA", "BROKER SOFTWARE", and "API REFERENCE". The main content area contains an introduction, a description of the broker's purpose, and a list of "Access Data" resources.

Docs » Pitt-Google Broker [Edit on GitHub](#)

Pitt-Google Broker

The Pitt-Google Broker is a cloud-based alert distribution service designed to provide near real-time processing of data from large-scale astronomical surveys like the [Legacy Survey of Space and Time](#) (LSST). LSST will deliver on order a million real-time alerts each night providing information on astronomical targets within 60 seconds of observation. The Pitt-Google Broker is a scalable broker system being designed to maximize the availability and usefulness of the LSST alert data by combining cloud-based analysis opportunities with value-added data products.

The Pitt-Google Broker runs on the [Google Cloud Platform](#) (GCP) and is currently focused on processing and serving alerts from the [Zwicky Transient Facility](#) (ZTF), and extending broker capabilities using ZTF, the LSST Alert Simulator, and the DECam Alliance for Transients (DECAT) stream.

Access Data

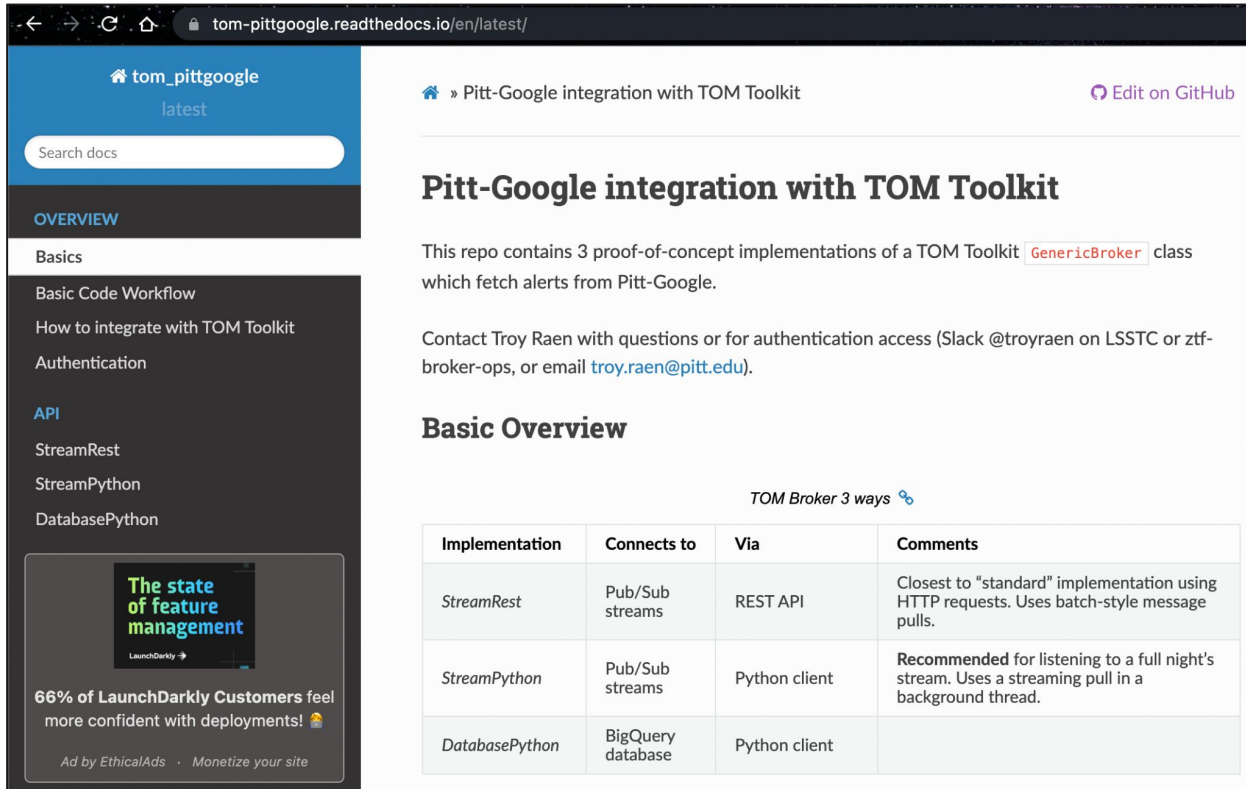
- [Data Overview](#)
 - [Pub/Sub](#)
 - [BigQuery](#)
 - [Cloud Storage](#)
 - [Note on Costs](#)
- [Initial Setup](#)
 - [Step 1: Create a GCP project](#)
 - [Step 2: Configure authentication](#)
 - [Step 3: Enable and install APIs](#)
 - [Cleanup: Delete a GCP project](#)
- [BigQuery Databases](#)
 - [Prerequisites](#)

Pitt-Google does not have to write its own APIs.

(And end-users do not have to rely on Pitt-Google to write and maintain a good API.)

We do, however, provide tutorials and Python wrappers of Google APIs for accessing Pitt-Google data resources.

Data Access from Google Cloud Services



The screenshot shows a GitHub repository page for 'Pitt-Google integration with TOM Toolkit'. The page includes a navigation sidebar on the left with sections for 'OVERVIEW', 'Basics', and 'API'. The main content area features a title, a description of the repository's purpose, contact information for Troy Raen, and a 'Basic Overview' section containing a table of implementation details.

Pitt-Google integration with TOM Toolkit

This repo contains 3 proof-of-concept implementations of a TOM Toolkit `GenericBroker` class which fetch alerts from Pitt-Google.

Contact Troy Raen with questions or for authentication access (Slack @troyraen on LSSTC or ztf-broker-ops, or email troy.raen@pitt.edu).

Basic Overview

TOM Broker 3 ways

Implementation	Connects to	Via	Comments
<code>StreamRest</code>	Pub/Sub streams	REST API	Closest to "standard" implementation using HTTP requests. Uses batch-style message pulls.
<code>StreamPython</code>	Pub/Sub streams	Python client	Recommended for listening to a full night's stream. Uses a streaming pull in a background thread.
<code>DatabasePython</code>	BigQuery database	Python client	

The state of feature management
LaunchDarkly

66% of LaunchDarkly Customers feel more confident with deployments! 🎉

Ad by EthicalAds · Monetize your site

REST and client APIs make integration with external services straightforward.

Google Cloud quotas and limits



- There are quotas and limits on how much you can use. Throughput, API calls, startup time, etc.
- They are broken down into a fine level of detail.
- These are generally high, but some may be an issue at LSST-scale.
- Some you can request to increase, some you can't.
- An example:
 - <https://cloud.google.com/bigquery/quotas>: Maximum number of concurrent interactive queries: 100

Google Cloud cost model

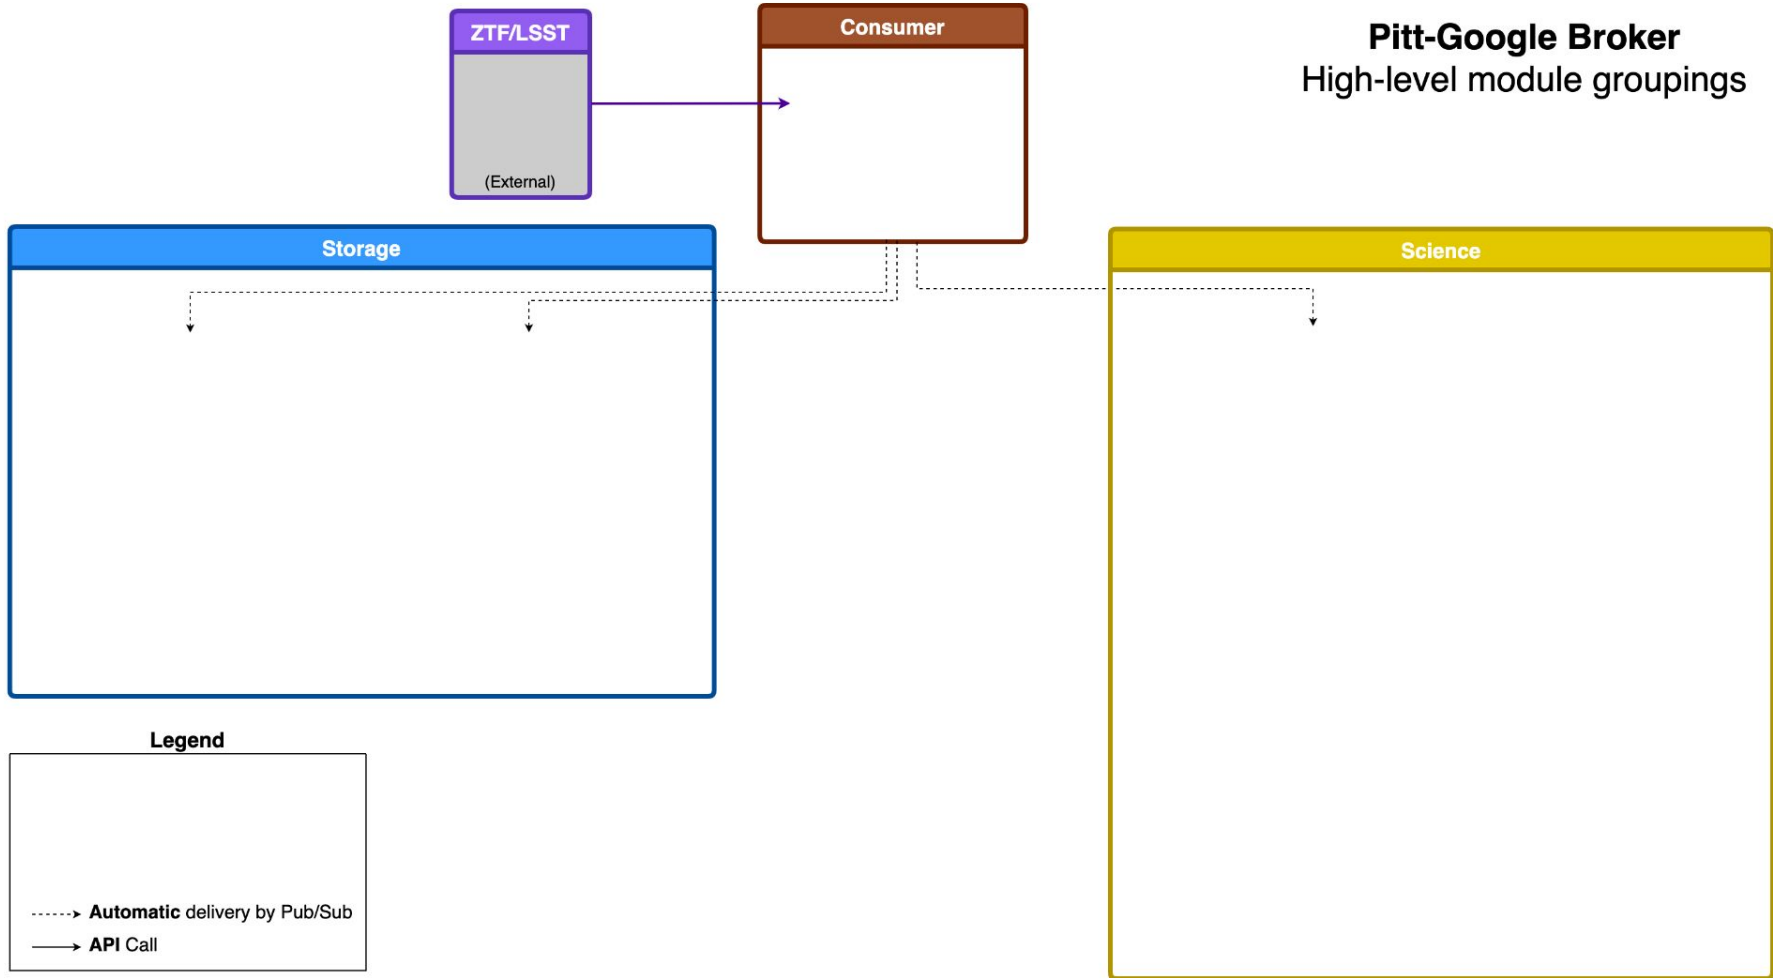


- **Free to connect, test things out, run ongoing analysis at low volume** (low from LSST perspective)
 - \$300 credit for new accounts
 - Free tier quotas renew monthly for everyone (e.g., BigQuery: The first 1 TB of query data processed per month is free.)
- **Beyond the free stuff, everyone pays for what they use.** Pay as you go. No upfront costs.
 - Example: We pay to store data in BigQuery, and query it ourselves. Our end-users pay for their own queries on data we have stored.
- **Costs broken down to a fine level of detail.** Some examples:
 - <https://cloud.google.com/bigquery/pricing>: The first 1 TB of query data processed per month is free, after which it's \$5.00 per TB.
 - <https://cloud.google.com/run/pricing>: First 180,000 vCPU-seconds free per month, after which it's \$0.00002160 / vCPU-second (Tier 2, Sao Paulo, Brazil).
 - Watch out for egress charges. Consider computing next to the data (in the Cloud).
- **Tradeoffs:**
 - Paying a cloud provider to manage infrastructure, run your jobs, and store your data.
 - Not paying to purchase/manage hardware; network bandwidth; personnel time to plan, develop, manage, and support most of the infrastructure.
 - Organizational budgets may be scoped for some things, but not others.

Pitt-Google Broker's Architecture and Google Cloud Services

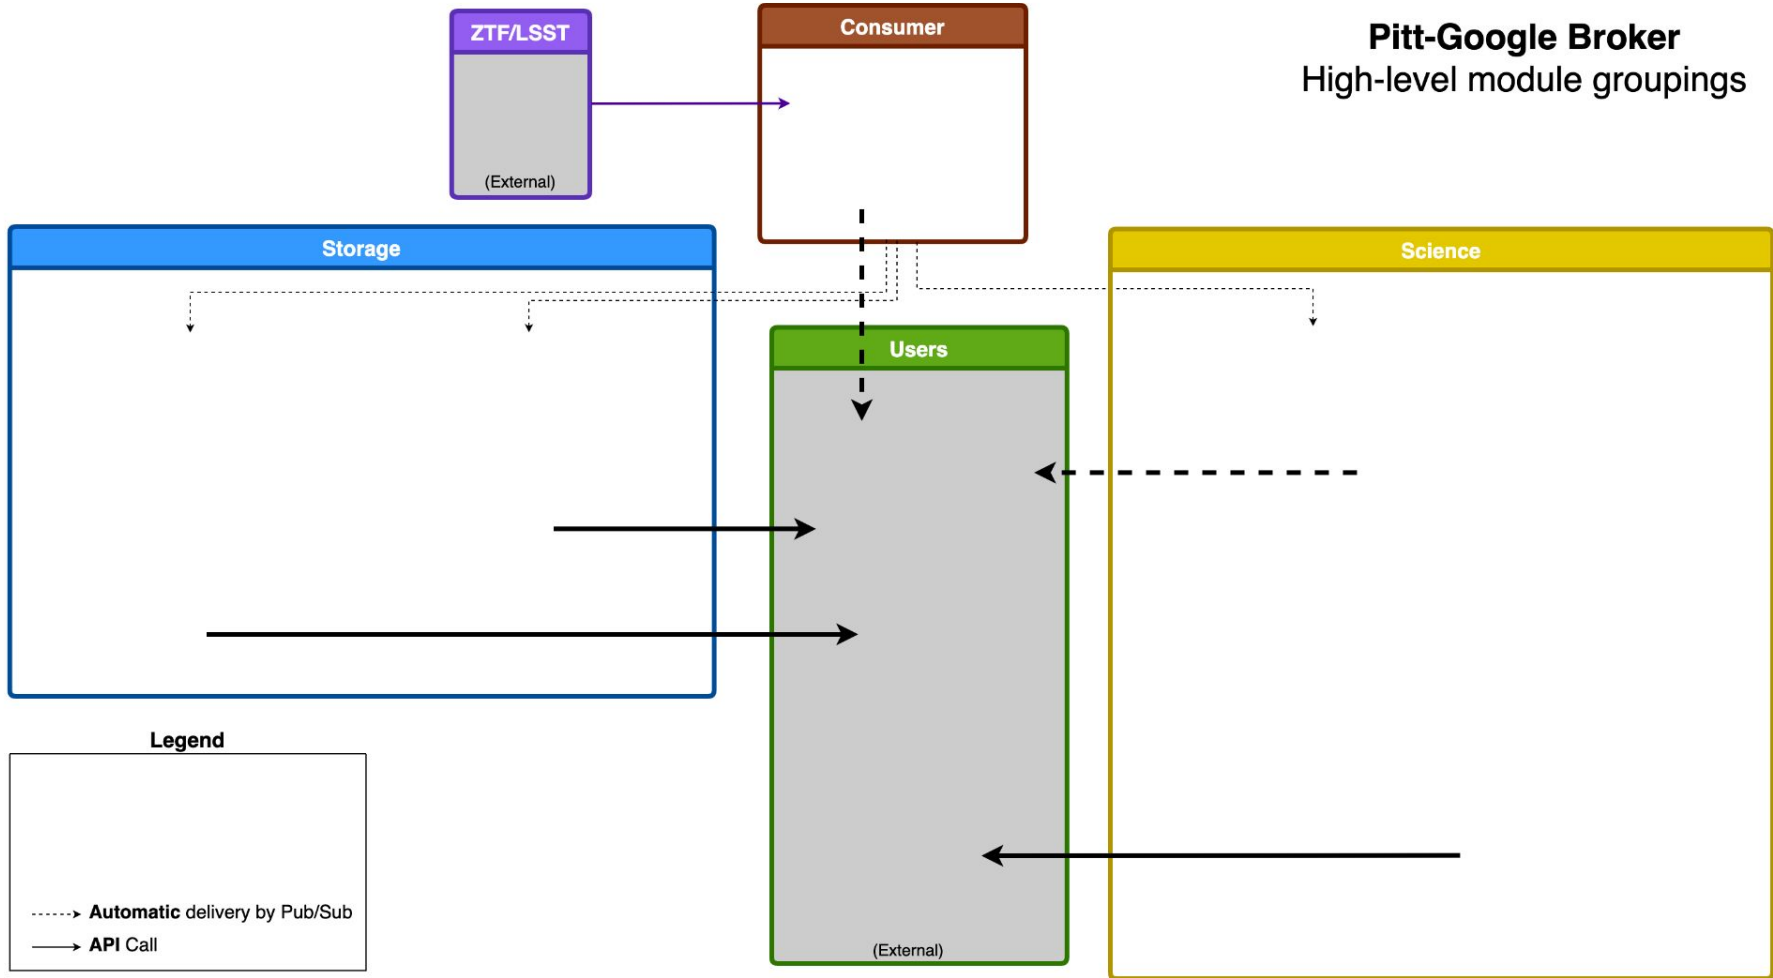
Pitt-Google Broker

High-level module groupings

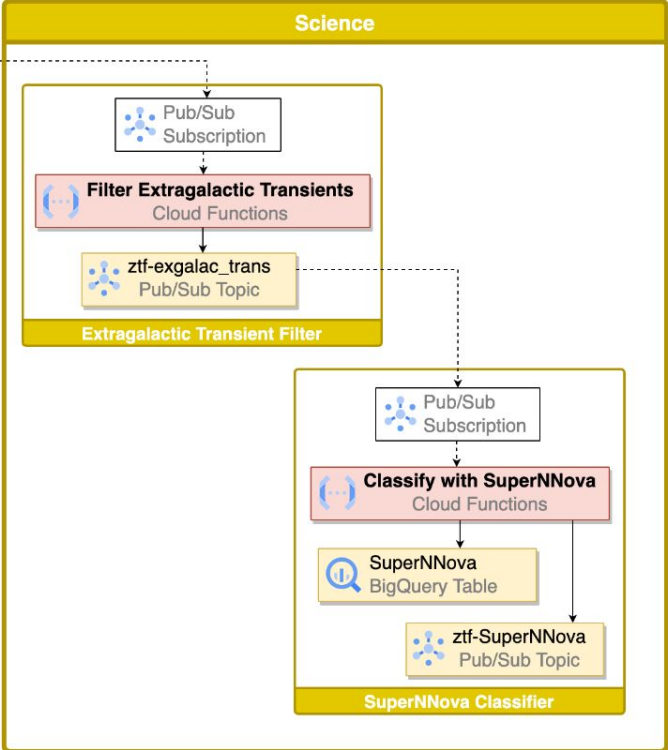
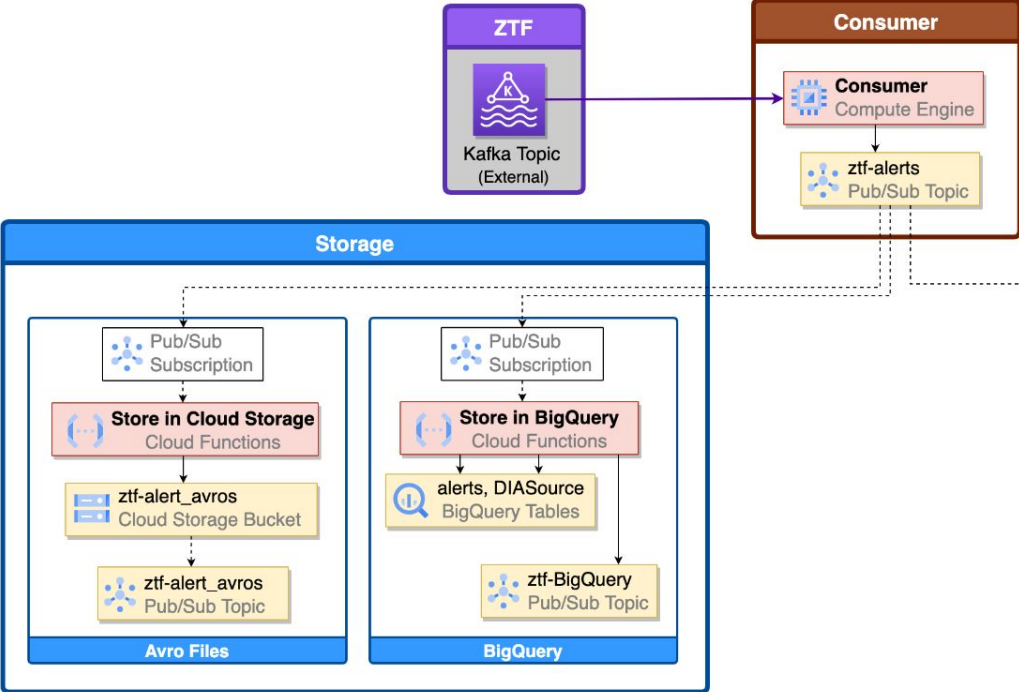


Pitt-Google Broker

High-level module groupings



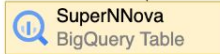

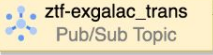


Pitt-Google Broker - Current



Legend

- Public Data Resource** (Yellow box)
- Private Data Resource** (White box)
- Compute Resource** (Pink box)
- Automatic delivery by Pub/Sub** (Dashed arrow)
- API Call** (Solid arrow)

Pitt-Google Broker - Current



Legend

- Public Data Resource** (Yellow box)
- Private Data Resource** (White box)
- Compute Resource** (Red box)
- Automatic delivery by Pub/Sub** (Dashed arrow)
- API Call** (Solid arrow)

Event-driven microservices model



Jargon: Microservices

“is an architectural style that *structures an application as a collection of services* that are ... *loosely coupled, independently deployable, ...*”

<https://microservices.io/>

“Events” are alerts or messages.

Per-message processing. No batching.

Benefits: Fast. Lightweight (no VM/cluster management). Per-alert processing easy to think about and develop. Fault-tolerant (failures automatically redelivered by Pub/Sub, other messages not affected). Scales (from zero) automatically and nearly instantaneously to meet demand.

Drawbacks: Cross matches are a challenge within the pipeline. Can't query an external service like CDS 10 million times per night. Per-alert lookups in BigQuery are expensive.

Pub/Sub



<https://cloud.google.com/pubsub/docs/overview>

- Pub/Sub is a streaming message service.
- Asynchronous publish-subscribe. Publishers are decoupled from subscribers.
- Messages retained in subscription for (max) 7 days, or until acknowledged by subscriber.
- **Push** to any HTTP endpoint for event-driven processing/storage.
 - Google Cloud Functions and Cloud Run
 - Dataflow
 - AWS Lambda (should be doable, but haven't tried it yet)
- **Pull** from anywhere with network access. ([APIs](#): Python, REST, gRPC, CLI, ...).
- Pricing basics:
 - Throughput: First 10 GiB/month is free. After that, \$40 /TiB.
 - Egress for throughput that crosses Google Cloud regions. \$0.08 - \$0.12 per GB in North and South America
 - No storage fees, unless you turn on special retention options.
 - <https://cloud.google.com/pubsub/pricing>

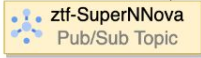

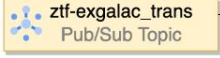





BigQuery



<https://cloud.google.com/bigquery/docs/introduction>

- BigQuery is a fully managed data warehouse. (can analyze external data sources)
- SQL queries (ANSI:2011 compliant)
- Columnar storage, optimized for complex analytical queries on large datasets
- Supports views & materialized views, and user-defined functions
- Streaming row inserts (data available to query within seconds), or batch load
- Full [API](#) access
- Built-in GIS, which we can exploit for positional matching with static catalogs... but...
- No indexes
 - Does a full column scan for every query. Single point lookups are expensive. Can reduce this by partitioning and clustering the tables
 - Cross matching within the event-driven pipeline is challenging. Currently we are working on adding a spatial index to AllWISE catalog and clustering. Then will test (HEALPix. Also explored HTM)
- <https://cloud.google.com/bigquery/pricing>: First TB of query data processed per month is free, then \$5.00 per TB. Active storage: first 10 GB/month is free, then \$0.02 per GB.

Pitt-Google Broker - Current



Legend

Public Data Resource

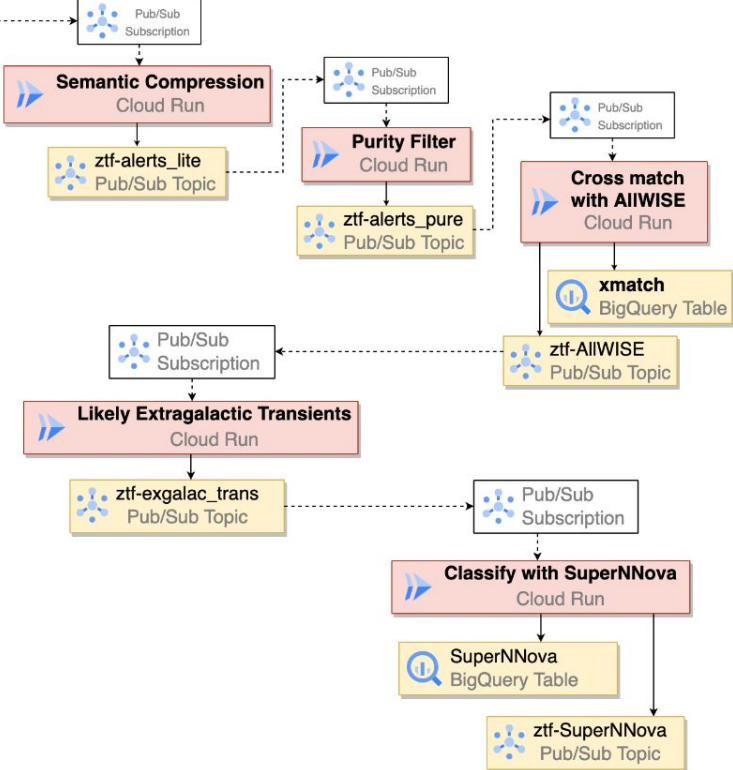



Private Data Resource

Compute Resource

-----> **Automatic** delivery by Pub/Sub

-----> **API Call**

Pitt-Google Broker - Near Future



Legend

- Public Data Resource
- Private Data Resource
- Compute Resource
- > Automatic delivery by Pub/Sub
- > API Call

Pitt-Google Broker - ~1 Year Future

ZTF/LSST

 Kafka Topic (External)

Consumer
 Compute Engine

ztf-alerts
 Pub/Sub Topic

Pub/Sub Subscription
Store in Cloud Storage
 Cloud Run

ztf-alert_avros
 Cloud Storage Bucket

ztf-alert_avros
 Pub/Sub Topic

Pub/Sub Subscription
Store in BigQuery
 Cloud Run

alerts, DIASource, DIAObject
 BigQuery Tables

ztf-BigQuery
 Pub/Sub Topic

User

mysubscription1
 Pub/Sub Subscription

myclassifier
 Cloud Run

mytable1
 mytable2
 mytable3
 BigQuery Tables

myanalysis
 Compute Engine

myfilter
 Cloud Functions

mysubscription2
 Pub/Sub Subscription (External)

Pub/Sub Subscription
Semantic Compression
 Cloud Run

alerts_lite
 Pub/Sub Topic

Pub/Sub Subscription
Purity Filter
 Cloud Run

alerts_pure
 Pub/Sub Topic

Pub/Sub Subscription
Cross match with - AllWISE - Gaia
 Cloud Run

xmatch
 BigQuery Table

xmatch
 Pub/Sub Topic

Pub/Sub Subscription

Classifiers - SuperNNova - Variable stars
 Cloud Run

Classifications
 BigQuery Table

classifications
 Pub/Sub Topic

Pub/Sub Subscription
Filters
 Cloud Run

Likely Extragalactic
 Likely Supernova
 Likely Variable Star
 Likely Cataclysmic Variable
 Pub/Sub Topics

Legend

Public Data Resource (Yellow box)

Private Data Resource (White box)

Compute Resource (Pink box)

-----> **Automatic delivery by Pub/Sub**

-----> **API Call**

Filters on Pub/Sub streams



Reason: Ease analysis by curating streams to deliver only the subset of alerts that are likely of interest to a particular researcher/working group.

Options:

1. Pitt-Google broker runs a filter producing a curated stream.
2. End user places a filter directly on a subscription.
 - a. Restrictions: Filter can only access data explicitly placed in message attributes, not the full alert packet . A limited set of operators is available (e.g.: = and !=, but not > or <)
3. End user listens to a stream and runs their own filter.
 - a. Simple to implement using Cloud Functions/Run

Pitt-Google will:

1. Produce multiple curated streams supporting research in LSST's major science goals.
2. Support filters directly on subscriptions by thinking carefully about which data to expose and how.
3. Enable users running their own filters by providing tutorials and guidance.

Cloud Functions and Cloud Run



Cloud Functions

<https://cloud.google.com/functions/docs>

Serverless, highly scalable, **functions** as a service.

Triggered by Pub/Sub messages (or HTTP requests). Event-driven processing.

Max runtime for single invocation: 9 minutes

No runtime environment management. Limited to (Python runtime available)

A single instance can process one request at a time. Number of instances automatically scales to handle all incoming requests.

Cloud Run

<https://cloud.google.com/run/docs>

Serverless, highly scalable, compute to run **containers**.

Triggered by requests. Event-driven processing.

Max runtime for single invocation: 60 minutes.

Any environment, language.

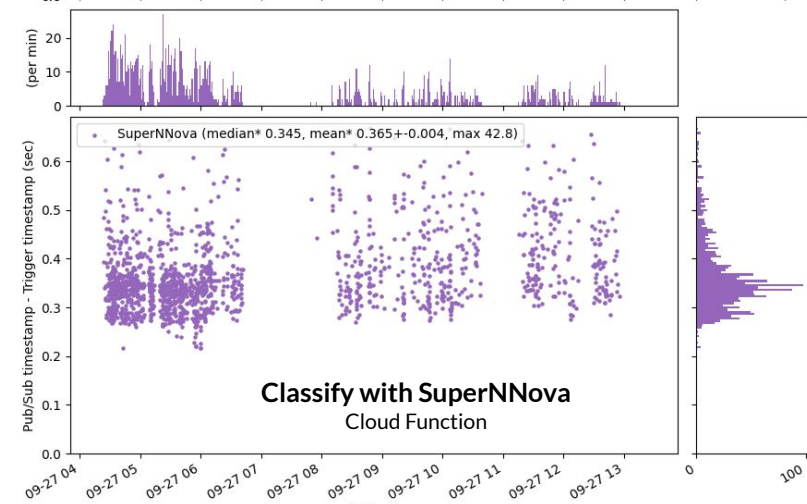
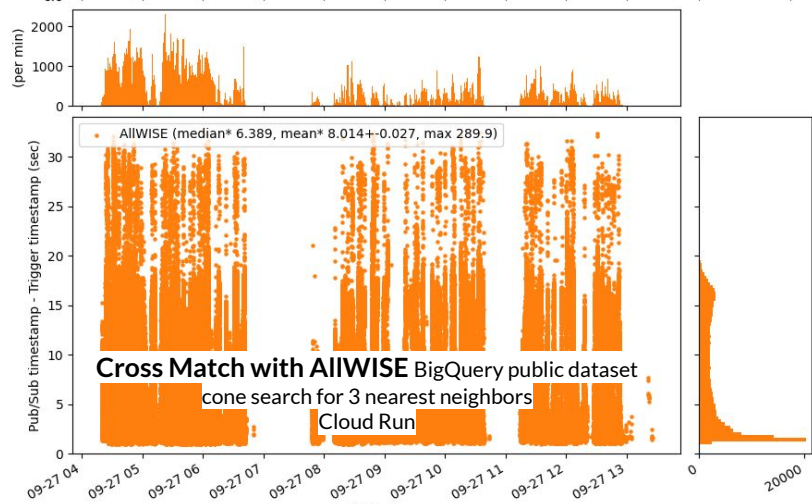
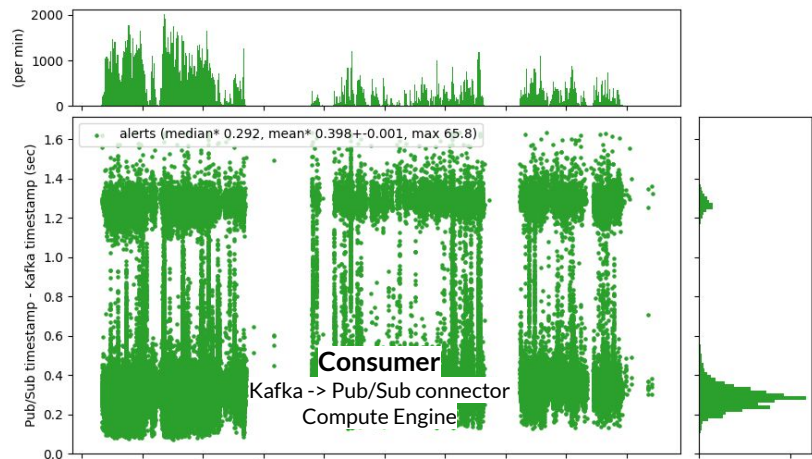
A single container instance can process up to 1000 concurrent requests. Number of instances automatically scales to handle all incoming requests.

Consumer: Kafka -> Pub/Sub Connector




- Single “g1-small” VM
 - 0.5 vCPU
 - 1.70 GB memory
 - ~\$8/month
- Running a Kafka Connect plugin
 - Passes alert bytes straight through to Pub/Sub, no decoding. Also passes message attributes.

Module processing times on ZTF stream



Module processing times at LSST rates (briefly)



Other Google Cloud Tools

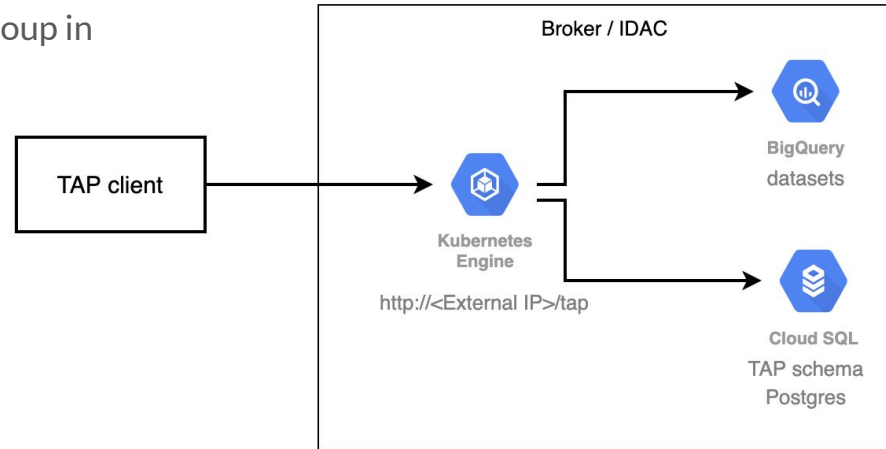
Things I'd like to explore.

Things that may be of use to the LineA IDAC.

BigQuery + TAP and ADQL

TAP and ADQL for BigQuery was announced in the 235th AAS Meeting by [J. Thomson et al. \(2020\)](#)

Currently under development by a private group in collaboration with Google.



International Virtual Observatory Alliance

[Table Access Protocol](#)

[Astronomical Data Query Language](#)

Used in the Rubin Science Platform

BigQuery ML



<https://cloud.google.com/bigquery-ml/docs/introduction>

“create and execute machine learning models in BigQuery using standard SQL queries”

Compute at the data.

BigQuery Public Dataset

London bicycle hires k-means example.

<https://cloud.google.com/bigquery-ml/docs/kmeans-tutorial>

Run the CREATE MODEL query

```
CREATE OR REPLACE MODEL
  bqml_tutorial.london_station_clusters OPTIONS(model_type='kmeans',
  num_clusters=4) AS
WITH
  hs AS (
  SELECT
    h.start_station_name AS station_name,
  IF
    (EXTRACT(DAYOFWEEK
      FROM
        h.start_date) = 1
    OR EXTRACT(DAYOFWEEK
      FROM
        h.start_date) = 7,
    "weekend",
    "weekday") AS isweekday,
  h.duration,
  ST_DISTANCE(ST_GEOGPOINT(s.longitude,
    s.latitude),
    ST_GEOGPOINT(-0.1,
      51.5))/1000 AS distance_from_city_center
  FROM
    `bigquery-public-data.london_bicycles.cycle_hire` AS h
```

BigQuery + Data Studio

<https://cloud.google.com/bigquery/docs/visualize-data-studio>

Pitt-Google
dataset.

Dwarf nova
lightcurve example.

The screenshot displays the Google BigQuery interface. The top navigation bar includes 'EDITOR', 'COMPOSE NEW QUERY', and 'UNSAVE... 2'. The main content area shows a query editor with the following SQL query:

```
1 SELECT objectId, candid, jd, fid, magpsf, sigmapsf FROM `ardent-cycling-24341`
```

Below the query editor, the 'Query results' section is visible, showing a table with 11 columns: Row, objectId, candid, jd, fid, magpsf, and signr. The table contains 4 rows of data. A tooltip for 'Explore with Data Studio' is overlaid on the results.

Row	objectId	candid	jd	fid	magpsf	signr
1	ZTF18abccnr	1586429820815015004	2459340.9298264	2	20.140676498413086	0.
2	ZTF18abccnr	1586429820815015004	2459340.9298264	2	20.140676498413086	0.
3	ZTF18abccnr	1602457910815015005	2459356.9579167	1	20.446569442749023	0.1
4	ZTF18abccnr	1602457910815015005	2459356.9579167	1	20.446569442749023	0.1


At the bottom of the interface, there are tabs for 'PERSONAL HISTORY', 'PROJECT HISTORY', and 'SAVED QUERIES'. The page number '45' is visible in the bottom right corner.

BigQuery + Data Studio

<https://cloud.google.com/bigquery/docs/visualize-data-studio>

Pitt-Google
dataset.

Dwarf nova
lightcurve example.



Summary

A selection of takeaways

Google Cloud Model:

- Services
 - Let Google handle hardware, scaling, fault tolerance, data storage & delivery... especially at LSST scales
 - You focus writing your data analysis
- BigQuery: Fully managed data warehouse. SQL access. Optimized for complex analytic queries at petabyte-scale.
- Costs:
 - Everyone pays for what they use.
 - Beware of egress; consider compute next to the data instead.
 - Budget scopes: Paying cloud provider for services, not humans in your organization for their time/effort.

Pitt-Google alert broker:

- Motivation:
 - Enable scientific collaboration by making all data available to everyone via well-developed and stable tools for data analysis/deliver/storage at this scale.
 - Support manageability via light filtering and value added data (cross match and classification).
- Taking advantage of Google Cloud services to simplify broker design and management.
- Event-driven data pipeline using Pub/Sub, Cloud Run/Functions, BigQuery.
 - Fast, fault tolerant, scales automatically, easy to build and manage.