

Rubin Science Platform on Google: the story so far.

William O’Mullane,¹ Frossie Economou,¹ Flora Huang,² Dan Speck,³
Hsin-Fang Chiang,¹ Melissa L. Graham,⁴ Russ Allbery,¹ Christine Banek,¹
Jonathan Sick,¹ Adam J. Thornton,¹ Jess Masciarelli,² Kian-Tat Lim,⁵
Fritz Mueller,⁵ Sergey Padolski,⁶ Tim Jenness,¹ K. Simon Krughoff,¹
Michelle Gower,⁷ Leanne P. Guy,¹ and Gregory P. Dubois-Felsmann,⁸

¹*Rubin Observatory Project Office, 950 N. Cherry Ave., Tucson, AZ 85719, USA*

²*Google Inc., 1600 Amphitheatre Parkway, Mountain View, CA 94043, USA*

³*Burwood Group, 125 South Wacker Drive Suite 2950, Chicago, IL 60606, USA*

⁴*University of Washington, Dept. of Astronomy, Box 351580, Seattle, WA 98195, USA*

⁵*SLAC National Accelerator Laboratory, 2575 Sand Hill Rd., Menlo Park, CA 94025, USA*

⁶*Brookhaven National Laboratory, Upton, NY 11973, USA*

⁷*NCSA, University of Illinois at Urbana-Champaign, 1205 W. Clark St., Urbana, IL 61801, USA*

⁸*IPAC, California Institute of Technology, MS 100-22, Pasadena, CA 91125, USA*

Abstract.

In 2020 Rubin Observatory made a tender for a three year interim data facility. This was won by Google. In partnership with Google we have now deployed the first externally used version of the Rubin Science Platform consisting of a Jupyter environment and a portal. Sitting behind this is TAP access to our project developed large scale database Qserv holding a synthetic catalogue of 140 million objects and around 500TB of simulated images. The initial users were given access to this in July 2021.

1. Introduction

The Legacy Survey of Space and Time (Ivezić et al. 2019) is an optical/near-IR survey of half the sky in ugrizy bands to $r \leq 27.5$ (36 nJy) based on 825 visits over a 10-year period: deep wide fast. Carried out by Rubin observatory on Cerro Pachón Chile, the survey will produce around 100 PB of data consisting of about a billion 16 Mpix images, enabling measurements for 40 billion objects!

We have recently begun a preoperations program which includes exposing users to the Science Platform.

2. Interim Data Facility on Google

We decided on an Interim Data Facility (IDF) to mitigate the risks of — and the delay imposed by — the US Data Facility(USDF) selection process. We felt this an excellent transition activity to prepare the operations team and Community for Rubin data. In a competitive tender we selected Google for this three year activity. We have installed construction delivered services (Rubin Science Platform and middleware/plumbing) on Google Cloud Platform. The architecture is depicted in Figure 1. We primarily leverage Google Kubernetes Engine (GKE) and Google Cloud Storage (GCS) to scale up and down as needed while batch processing images and serving the data release products. The full stack was automated using Terraform with help from the Burwood Group, which is not highlighted in the diagram but is outlined in section 2.1. In addition running services at scale improves readiness for Data Production activity during full survey operations.

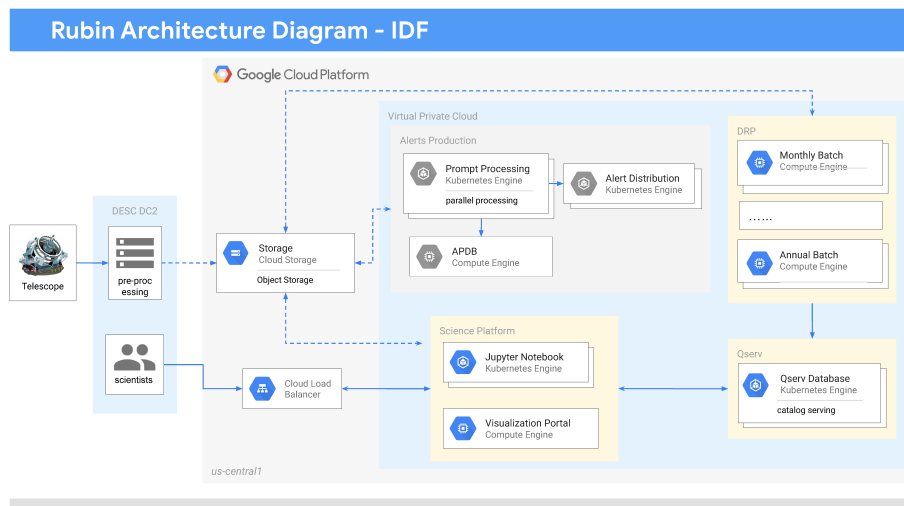


Figure 1. Architecture of the Interim Data Facility on Google, gray Items (alerts) are not yet implemented but shall be in 2022.

2.1. Docker, Kubernetes, Terraform

The setup of the infrastructure such as clusters on GCP is done using Terraform¹. Typically we have development, integration and production clusters for each major deployment e.g. Science Platform and Qserv. The orchestration of the Rubin application containers is done using Kubernetes² with deployment handled by ArgoCD³. Our build system, based on Jenkins and GitHub Actions, produces Docker containers for deployment (see e.g., Jenness et al. 2018).

¹<https://www.terraform.io/>

²<https://kubernetes.io/>

³<https://argoproj.github.io/cd/>

3. Data Previews

The IDF will be used for a series of Data Previews (DP; Marshall 2020). DP0 is using simulated data, DP1 commissioning camera data, and DP2 LSSTCam commissioning data. DP0 currently has two phases (O’Mullane 2020) with one focussing on data access and the other on data processing. In June 2021 data.lsst.cloud was opened up to the first 200 delegates as part of Data Preview 0.1 with data from the DESC DC2 simulated data (Abolfathi et al. 2021). For Data Preview 0.2 we shall reprocess the DESC data (DP0.1) with the LSST Science Pipelines v23, to generate a fully self-consistent data release. This will entail producing all the data products and catalogs as well as loading the catalogs in Qserv. Thus demonstrating a portable set of cloud enabled tools based on Butler Gen3 (Jenness et al. 2019) and PanDA.

4. Community engagement

The Rubin Community Engagement Team (CET) provided documentation and support to the scientists and students participating in DP0 as delegates. The term “DP0 delegates” was adopted to reflect DP0 participants’ important roles of representing the broad science community as learners, testers, and providers of feedback, and of sharing the benefits of their DP0 participation with their communities as teachers and colleagues.

Documentation at dp0-1.lsst.io included a delegate’s homepage listing all the resources, events, guidelines, and expectations; descriptions of the data products and table schema; instructions for accessing the RSP’s three aspects and usage risks; and tutorials to demonstrate simple science investigations with the DP0 data set at the RSP. Delegate resources included biweekly virtual “Delegate Assemblies” with hands-on demonstrations and breakout rooms to facilitate co-working; a dedicated category in the Rubin Community Forum (<https://community.lsst.org>) where delegates can ask questions and discuss their DP0 work; and a GitHub repository where delegates can contribute to shared code and notebooks. GitHub Issues were also used as a means for delegates to request technical assistance.

Rubin staff from the CET and Data Production worked together to respond to delegate’s requests for support and bug reports. All of the delegate resources are publicly accessible to non-delegates, including the assemblies (and their recordings) and the Community Forum discussions. Some aspects of the Community Engagement model as applied to DP0 are designed to scale to thousands of users for the future Data Previews, such as the extensive documentation, the tutorials, and the use of the Community Forum. However other aspects, such as the frequent live virtual sessions, would not scale well to thousands of users and were instead specifically designed to ramp up skill levels from novice to intermediate in order to seed expertise with Rubin software in the community. Looking forward to future Data Previews, the Community Engagement strategy will focus on building and fostering a vibrant science community, supported by infrastructure, that is able to help itself and to crowd-source solutions.

5. Conclusion and lessons learned

User feedback has been very positive on this initial Science Platform. Users do need a lot of guidance if we are to be efficient. For example Qserv, accepting ADQL, can

build great histograms (minutes for full dataset) but users tended to do this in Python by running a query to extract all the data and then binning it in Python. AutoScale really works – but it's not really fast considering our containers, it takes minutes to spin up new nodes as far as users concerned! DASK looks very interesting but we are not sure how to manage resource allocations yet. No matter how much you prepare you always learn something new! The Data Butler abstracts data access from the user and we designed the Butler to support pluggable backends from the beginning. Deploying the Butler to Google was relatively straightforward since Butler already supported Postgres for the registry and S3 for file access.

Management buy-in is a big part of any cloud adventure, there is a lot of reluctance to change from in-house models which requires much convincing. One thing we did were proof of concept studies with both Amazon (Bektsevich et al. 2020; Chiang et al. 2020) and Google (Lim 2019). Each ran for six months and demonstrated we could deploy and operate on either platform. This allayed concerns about vendor lockin and improved cost models.

Since we built cloud ready services, setup on GCP was fairly painless and smooth. Support from Google and Burwood was excellent and rapid, and weekly tag-ups continue to make sure everything is in place or planned. Since Rubin is between data facilities for some years this is a quite cost effective approach.

References

- Abolfathi, B., et al. 2021, *The Astrophysical Journal Supplement Series*, 253, 31. URL <http://dx.doi.org/10.3847/1538-4365/abd62c>
- Bektsevich, D., Chiang, H.-F., Lim, K.-T., Miller, T. L., Thain, G., Jenness, T., Bosch, J., Salnikov, A., & Connolly, A. 2020, arXiv e-prints, arXiv:2011.06044. 2011.06044
- Chiang, H.-F., Bektsevich, D., & the AWS-PoC team 2020, AWS Proof of Concept Project Report. URL <http://DMTN-137.lsst.io>
- Ivezić, Ž., et al. 2019, *ApJ*, 873, 111. 0805.2366
- Jenness, T., Bosch, J., Schellart, P., Lim, K.-T., Salnikov, A., & Gower, M. 2019, in *Astronomical Data Analysis Software and Systems XXVII*, edited by P. J. Teuben, M. W. Pound, B. A. Thomas, & E. M. Warner, vol. 523 of *Astronomical Society of the Pacific Conference Series*, 653. 1812.08085
- Jenness, T., et al. 2018, in *Software and Cyberinfrastructure for Astronomy V*, vol. 10707 of *Proc. SPIE*, 1070709
- Lim, K.-T. 2019, Google Cloud Engagement Results. URL <http://dmtn-125.lsst.io>
- Marshall, P. 2020, Release Scenarios for LSST Data. URL <https://docushare.lsstcorp.org/docushare/dsweb/Get/RDO-11>
- O'Mullane, W. 2020, Data Preview 0: Definition and Planning. URL <http://RTN-001.lsst.io>