

When clouds might be good for LSST

William O'Mullane

2019-10-23

1 Introduction

The Large Synoptic Survey Telescope (Large Synoptic Survey Telescope (LSST)) will usher in a new era of data-intensive astronomy. The observing program will observe the southern sky repeatedly over 10 years in 6 bands providing an unprecedented census of the astrophysical bodies in the universe. Funded by the National Science Foundation (NSF) and the Department of Energy (DOE), this keystone observatory is due to go into operations in October 2022.

Producing 20TB of data a night, this is a huge step up in data acquisition from other optical telescopes. At its conception this was considered an ominous data volume requiring highly specialized computing infrastructure. In the intervening time, however, the growth of planetary-scale industry services (such as Google or Facebook) has resulted in software, engineering techniques and infrastructure that render this sort of data flow routine. LSST operations are expected to cost tens of millions a year with an order of \$10M computing budget.

The LSST computing load is a poster child for cloud computing - the science platform is designed for Kubernetes and the data release processing fits perfectly with opportunistic compute pricing. This is a large NSF project and the one most suited and ready for cloud deployment. We have used some pathfinder deployments to demonstrate that it is feasible to use commercial cloud providers for the LSST Data Management production system. Such a move would bring significant technological and operational advantages; the barrier to acting on this is price and uncertainty on future pricing.

A solution might be to reach a fixed price partnership for a cloud-based deployment of the Data Management systems in which Google undertake to provide *do what is needed* for success at some reasonable and agreed-upon annual fee.

2 Studies to date

We have demonstrated some of the major components of LSST Data Management work well on Google [DMTN-072].

We deployed Qserv on Google with reasonable performance. We demonstrated Data transfer could adequate for Prompt Processing, within the limits of the available network. The Prompt Product Database was stood up and tested. The Science Platform was deployed and users simulated. The later of course is designed around kubernetes and made for Google Cloud.

A detailed report may be found in DMTN-125.

We have not considered bulk transfer of data to other partners - Google may be much better placed to do this than any computing centre.

3 LSST compute and storage needs

The greatest cost driver is storage - we accumulate about 50PB a year of data. All of this needs to be processed annually. Hence in year 10 we need to access about half an Exabyte of data. Not all of this will be regularly accessed, it is likely few of the raw images will be reprocessed by individual astronomers.

Table 1 gives a rough overview of compute and storage needs.

Table 1: Various inputs for deriving costs

Year	2019	2020	2021	2022	2023
FLOPs Needed Total (no Alerts)	1.00E+19	9.48261E+19	1.00E+19	4.74131E+20	5.91525E+20
Time to Process days	252.0	252.0	252.0	252.0	252.0
Time to Process seconds	21772800.0	21772800.0	21772800.0	21772800.0	21772800.0
Instantaneous GFLOP/ s	4.59E+02	4355.255691	4.59E+02	21776.27846	27168.07327
Instantaneous GFLOP/ s (inc Alerts)	4.59E+02	30025.25569	2.61E+04	21776.27846	27168.07327
Disk Space TB	5000	10000	20000	50000	100000
I/ O for year TB	15000	30000	60000	150000	300000
Base numbers	GFLOP				
LDM-138 DR1,2 Data Rel sheet row 1	426717500000		97381399021		
LDM-138 DR3 Data Rel sheet row 2	959090000000				
LDM-138 Alert Instantaneous	25670				
Alert Total, assuming 275k visits/ year	177219625000				
Total Yr1 (inc DAC)	474130555556				

4 Cost

Currently the cloud computing cost models do not align well with federal research computing plans. We believe this gap will narrow or disappear in five years or so, however in the meantime LSST and Google may miss an opportunity to move a major research project on to commodity cloud computing.

Using the information from the Google study where we ran some of our real processes, we have come to a price for running the Science Platform and storing data on Google. For 15PB of storage and a modest K8S cluster to host the platform the projected 2022 cost is around \$3M of which more than $\frac{2}{3}$ are storage costs.

Though a good price this is not a sustainable price for LSST, we can construct petascale storage we would own and use for 5 years for under \$200K a petabyte (the implied *annual* price at google). We require about 50 Petabytes a year for 10 years. The out year costs look prohibitive on the cloud.

The cost of compute is probably not an issue in comparison - we can use spot/interruptable instance pricing for Data Release Production (DRP).

5 Conclusion

LSST Data Management (DM) has been constructing a cloud ready system for many years. We believe commercial cloud is the correct approach but we may be a few years ahead of commercial and federal cost models aligning. We hope that we may be able to partner with google to usher in a new era of federally funded research in the cloud.

References

- [1] [DMTN-125], Lim, K.T., 2019, *Google Cloud Engagement Results*, DMTN-125, URL <http://dmtn-125.lsst.io>

[2] [DMTN-072], O'Mullane, W., Swinbank, J., 2018, *Cloud technical assesment*, DMTN-072, URL <https://dmtn-072.lsst.io>, LSST Data Management Technical Note

A Glossary

Data Management The LSST Subsystem responsible for the Data Management System (DMS), which will capture, store, catalog, and serve the LSST dataset to the scientific community and public. The DM team is responsible for the DMS architecture, applications, middleware, infrastructure, algorithms, and Observatory Network Design. DM is a distributed team working at LSST and partner institutions, with the DM Subsystem Manager located at LSST headquarters in Tucson.

Department of Energy cabinet department of the United States federal government; the DOE has assumed technical and financial responsibility for providing the LSST camera. The DOE's responsibilities are executed by a collaboration led by SLAC National Accelerator Laboratory.

DM Data Management.

DOE Department of Energy.

DRP Data Release Production.

LSST Large Synoptic Survey Telescope.

National Science Foundation primary federal agency supporting research in all fields of fundamental science and engineering; NSF selects and funds projects through competitive, merit-based review.

NSF National Science Foundation.

Prompt Processing The processing that occurs at the Archive Center on the nightly stream of raw images coming from the telescope, including Difference Imaging Analysis, Alert Production, and the Moving Object Processing System. This processing generates Prompt Data Products.

Qserv Proprietary Database built by SLAC for LSST.

Science Platform A set of integrated web applications and services deployed at the LSST Data Access Centers (DACs) through which the scientific community will access, visualize, and perform next-to-the-data analysis of the LSST data products.