

# The Rubin Observatory's Legacy Survey of Space and Time DP0.2 processing campaign at CC-IN2P3

Quentin Le Boulch<sup>1,\*</sup>, Fabio Hernandez<sup>1</sup>, and Gabriele Mainetti<sup>1</sup>

<sup>1</sup>CNRS, CC-IN2P3, 21 avenue Pierre de Coubertin, CS70202, 69627 Villeurbanne CEDEX, France

**Abstract.** The Vera C. Rubin Observatory, currently in construction in Chile, will start performing the Legacy Survey of Space and Time (LSST) in 2025 for 10 years. Its 8.4-meter telescope will survey the southern sky in less than 4 nights in six optical bands, and repeatedly generate about 2 000 exposures per night, corresponding to a data volume of about 20 TiB every night. Three data facilities are preparing to contribute to the production of the annual data releases: the US Data Facility will process 35% of the raw data, the UK data facility will process 25% of the raw data and the French data facility, operated by CC-IN2P3, will locally process the remaining 40% of the raw data.

In the context of the Data Preview 0.2 (DP0.2), the Data Release Production pipelines have been executed on the DC-2 simulated dataset (generated by the Dark Energy Science Collaboration, DESC). This dataset includes 20 000 simulated exposures, representing 300 square degrees of Rubin images with a typical depth of 5 years.

DP0.2 ran at the Interim Data Facility (based on Google cloud), and the full exercise was independently replicated at CC-IN2P3. During this exercise, 3 PiB of data and more than 200 million files were produced. In this contribution we will present a detailed description of the system that we set up to perform this processing campaign using CC-IN2P3's computing and storage infrastructure. Several topics will be addressed: workflow generation and execution, batch job submission, memory and I/O requirements, etc. We will focus on the issues that arose during this campaign and how we addressed them and will present some perspectives after this exercise.

## 1 Introduction

The Vera C. Rubin Observatory is in the final stages of its construction on the Chilean Andes, and is preparing to execute the Legacy Survey of Space and Time (LSST) that will capture the entire southern sky every four nights in six different bands during ten years. The 8.4-meter telescope together with the 3.2 gigapixels camera (the largest digital camera ever built) will produce tens of petabytes of raw and calibration image data. Data processing pipelines will progressively generate an astronomical catalog of 20 billion galaxies and 17 billion stars and their associated physical properties that will allow scientists to explore the nature of dark matter and dark energy, study the Solar System and the Milky Way, as well as transient objects [1].

---

\*e-mail: [quentin.leboulch@cc.in2p3.fr](mailto:quentin.leboulch@cc.in2p3.fr)

This article presents the Data Preview 0.2 processing campaign performed at CC-IN2P3<sup>1</sup> for early integration tests of the Rubin data processing pipelines. Section 2 introduces the Rubin data management system and the processing pipelines, section 3 details the processing campaign that has been executed, and section 4 gives some perspectives for future improvements.

## 2 The Rubin Data Management System

The Rubin data management system is organized in three main components:

- the computing, storage and network infrastructure
- the middleware that provides data access and distributed processing
- the data processing pipelines

In this section we present the middleware components and the data processing pipelines.

### 2.1 The Data Butler

The Data Butler [2] is the software system that abstracts access to data. Individual sets of data (images and catalog data) are called datasets and are accessible in memory as Python objects through a Python API (a command line interface is also available), without the need to know where the files are located or what file format is used.

The datasets are organized in a database called the Registry which associates them with their astronomical description (sky coordinates, exposure identification, etc.). The Datastore is the component responsible for the storage of the files associated with the datasets, and provides the read and write interface to Python.

### 2.2 The processing pipelines

Data Release Production will be run every year in order to reprocess the complete set of observations from the beginning of the survey. The following processing operations are performed [3]:

- Characterization and calibration of individual exposures,
- Coaddition of single exposures to produce deeper images,
- Subtraction of images to identify time-variable sources,
- Processing of the coadded images for object detection and measurement,
- Measurements on single visit images,
- Catalog level processing to generate additional object properties.

Data processing pipelines [4] are composed of about 80 distinct processing tasks, implemented in Python with a C++ core. These tasks are connected through their input and output datasets (mostly catalog and image data), and the resulting Data Release Production processing workflow is represented as a directed acyclic graph (DAG). This large and complex workflow can be split into several sub-workflows called "steps" that correspond to a given level of processing that must be executed in a specific order. Each of these steps is composed of a number of pipeline tasks.


---

<sup>1</sup><https://cc.in2p3.fr>

### 2.3 The execution framework

The pipeline execution framework is called pipetask. It takes as input a description of the pipeline to be executed (list of tasks with their dependencies and configuration), the Data Butler repository where the input data are accessible, and the selection of data to be processed. pipetask provides the ability to execute the corresponding tasks in the correct order, and supports multiprocessing in case that multiple tasks can be executed in parallel when multiple CPU cores are available. It is however restricted to a single node, which is not sufficient for large scale workflows.

The Batch Processing Service (BPS) [5] is the system for distributed pipeline execution. It provides an additional layer that is able to execute multiple pipetask commands on independent nodes. BPS takes the same input information as pipetask and generates a workflow graph that contains all tasks: commands to be executed, resources requirements, and their dependencies. This workflow graph is then submitted to an external Workflow Management System (WMS) through a dedicated plugin for execution on a parallel computing infrastructure, as shown in figure 1.



**Figure 1.** Submission of a BPS workflow to a WMS for execution on a parallel computing infrastructure.

## 3 The Data Preview 0.2 processing campaign

### 3.1 Data Preview 0.2

Data Preview exercises aim at preparing Rubin systems before operations, and help the scientific community to get familiar with Rubin data. In particular, Data Preview 0 [6] uses simulated data to perform early integration tests of the processing pipelines and the Rubin Science Platform (RSP) [7] that provides viewing and analysis tools for images and catalogs. Catalog queries in the RSP are facilitated by Rubin’s high performance distributed database Qserv. Data Preview 0 also provides data for preparation of science analysis.

For Data Preview 0.2 (DP0.2) [8], simulated data provided by the Dark Energy Science Collaboration was reprocessed. These images were generated for their Data Challenge 2 [9] and represent 5 years of the LSST wide survey on 300 square degrees. This is approximately 0.5% of the full LSST survey that will cover 18 000 square degrees during 10 years. The DESC DC2 data set includes about 20 000 Rubin camera simulated exposures, each with up

to 189 detectors, which makes about 3 million files in total. The total volume of this data set is 50 TiB.

DP0.2 processing was executed at the Interim Data Facility, and it was decided to independently perform the same exercise at the IN2P3 / CNRS Computing Centre (CC-IN2P3) which hosts the French Data Facility. From the processing point of view, the purpose of this exercise was to test the latest versions of both the pipelines and the middleware, check the scalability of the infrastructure, and introduce automation of workflow execution.

After completion of the processing of these simulated data, the resulting data products were made available through the Rubin Science Platform.

### 3.2 Workflow execution

The WMS used at CC-IN2P3 for DP0.2 is Parsl [10], a tool that provides parallelism of Python programs on distributed computing resource. It was selected for its scalability, its ability to submit jobs to Slurm [11] which is the batch system used at CC-IN2P3, and its HighThroughputExecutor features which is suitable for CC-IN2P3 computing infrastructure. In addition, a Parsl plugin [12] for BPS was already available and only slight customisations were needed.


The full DP0.2 workflow is too large to be executed at once, so the execution was divided to be executed in several smaller submissions. First, as described in section 2.2, the workflow is divided into several distinct steps that are executed sequentially. In addition, for each step the input data was split into smaller independent groups. The way these groups have been defined takes into account the level of data on which the steps are executed (single exposures, or other partitions of the sky), as well as the resulting size of the generated workflow. On one hand, large workflows require a larger amount of resources to be generated by BPS and longer time to complete. On the other hand, small workflows imply a larger amount of independent submissions to be launched and monitored. Indeed, for each workflow execution it is necessary to identify potential errors during the processing and launch additional submissions to re-execute the missing parts. Finally, the clustering feature of BPS was also used to gather some tasks together in order to reduce the number of tasks to be executed and to improve scalability. The corresponding configuration files are stored in a GitHub repository [13].

### 3.3 Execution of the processing campaign

A Data Butler repository for DP0.2 was created and configured, and the input datasets were ingested [14]. The DP0.2 exercise was executed using BPS with the Parsl plugin. The workflow tasks ran within Parsl jobs that submitted to our Slurm computing platform. Before executing the main processing campaign, smaller tests were executed in order to check the Butler repository, the software environment, and progressively increase the number of parallel jobs. During the main campaign, up to 3 000 simultaneous Slurm jobs have been used, and more than 57 million tasks were executed, using more than 2 million of computing hours.

The tasks were very heterogeneous regarding the type of data they process and produce. The number of times each particular task must be executed varies a lot. In addition the memory and computing resources usage depends a lot on the type of task which is executed, and was not precisely known before launching the processing. Memory usage in particular can vary significantly for a given task with different input data (see figure 2). It was therefore necessary to continuously monitor the memory usage of the tasks and adjust the jobs configuration accordingly.

Using a custom wrapper script, the tasks were executed within Apptainer [15] containers, allowing us to isolate the task environment from the host's. Container images containing



**Figure 2.** Distribution of memory usage for the 7,344 deblend task executions during the DP0.2 processing campaign at CC-IN2P3.

the full Rubin software environment including the pipeline software were deployed using CernVM-FS [16].

The Butler repository Datastore was located on our CephFS [17] shared filesystem. 3 PiB of data and 201 M files have been generated during the exercise. The filesystem demonstrated sufficient performance in most of the cases. For one specific type of task we encountered I/O issues and the input files were copied on the local disk of the compute node to improve performance.

We used a PostgreSQL [18] database for the Butler repository registry. The database was 314 GiB in size at the end of the processing. In order to reduce the number of concurrent access, the BPS execution Butler feature [19] was used: jobs did not access directly the central database but used a local, smaller SQLite [20] database instead, where all necessary information was extracted from the Registry beforehand. At the end of the execution, information concerning the new files was ingested into the main Registry.


### 3.4 Analysis of the resources usage

An extensive analysis of the processing tasks resource usage is crucial in order to configure properly the Slurm jobs, but also to optimise the pipeline software, and correctly size our infrastructure.

The pipetask framework embeds its own profiling tool that measures the CPU and memory usage of each executed task in the pipeline. These metrics are stored as YAML files in the Butler repository. We retrieved these metrics for all tasks and exported them into CSV files that were analyzed to provide useful graphs.


Figure 3 shows the walltime and CPU time spent by all type of pipeline tasks. Note that some types of tasks imply multiple execution of the same task on different sets of input data. One can see that 90% of the computing time is spent on the 9 most consuming tasks. In addition, the CPU efficiency<sup>2</sup> for these tasks is rather good, which shows the good performances of our storage system.

<sup>2</sup>Understood as the fraction of elapsed time the CPU is actually doing computations, as opposed as waiting for data to process.



**Figure 3.** Elapsed and CPU time spent for each type of task. Items in the horizontal axis represent the various type of tasks executed in the campaign required to produce a data release, sorted by their aggregated elapsed time.

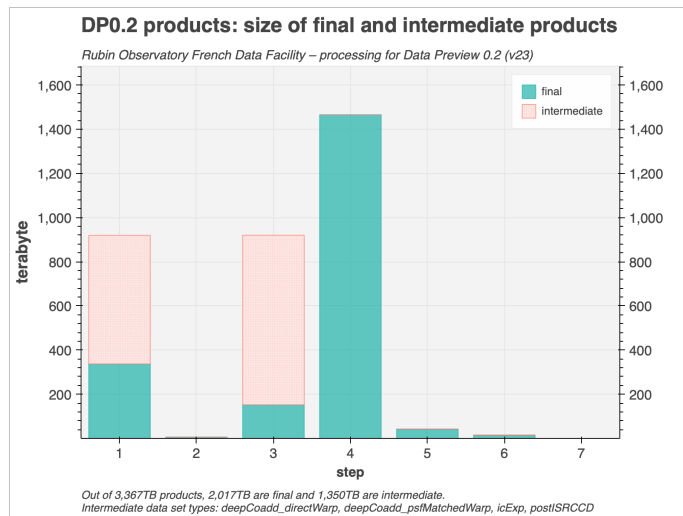
Figure 4 shows the distribution of the maximum memory usage per type of task. While most of the tasks have relatively low memory usage in the 1 to 10 GiB range, some tasks require several dozen of GiB to run and a few more than 100 GiB.



**Figure 4.** Memory consumption for each type of task. Whiskers show the 5th and 95th percentile.

We also collected and analyzed metrics regarding the files and data volume generated by the processing tasks. Figure 5 shows the total volume of data products generated by each of the 7 processing steps. Datasets have been tagged as final or intermediate. More than 1

PiB of data generated in step 1 and step 3 are intermediate files that can be deleted after the processing campaign is complete, however step 4 is the most data volume consuming step and generated more than 1.4 PiB of final products.



**Figure 5.** Total volume of final and intermediate data products generated by each DP0.2 processing step.

## 4 Summary and perspectives

The Vera C. Rubin Observatory is preparing to perform the Legacy Survey of Space and Time. 20 TiB of raw data will be generated every night for 10 years and the French Data Facility, operated by CC-IN2P3, will contribute resources for processing 40% of the cumulated data. In order to prepare our infrastructure we exercised the Data Release Production pipelines on the DC-2 simulated dataset composed of 20 000 simulated exposures.

This ambitious exercise was successful and showed that the French Data Facility was already able to execute the Rubin data processing pipelines at a significant scale. A similar exercise is ongoing using the dCache [21] storage system that would allow for a decision on what storage system to use for the operations phase.

One should also note that this processing campaign has been executed fully locally, however multi-site processing is being tested at smaller scale with the PanDA [22] system. It is also planned to perform more frequent tests with more recent release of the pipelines and of the execution environment. Finally, analysis of the resource usage will be used to understand how the CPU and memory usage of the pipeline can be improved.

## References

- [1] Z. Ivezić et al., *Astrophys. J.* **873**, 111 (2019)
- [2] T. Jenness, J.F. Bosch, A. Salnikov, N.B. Lust, N.M. Pease, M. Gower, M. Kowalik, G.P. Dubois-Felsmann, F. Mueller, P. Schellart, *The Vera C. Rubin Observatory Data Butler and pipeline execution system*, in *"Software and Cyberinfrastructure for Astronomy VII"* (2022), Vol. 12189 of *Proc. SPIE*, p. 1218911, arXiv:2206.14941

- [3] J. Swinbank, T. Axelrod, A. Becker, J. Becla, E. Bellm, J. Bosch, H. Chiang, D. Ciardi, A. Connolly, G. Dubois-Felsmann et al., *LDM-151 - Data Management Science Pipelines Design* (2020), Vera C. Rubin Observatory Data Management Controlled Document, <https://ldm-151.lsst.io/>
- [4] J. Bosch et al., *An Overview of the LSST Image Processing Pipelines*, in *Astronomical Data Analysis Software and Systems XXVII*, edited by P.J. Teuben, M.W. Pound, B.A. Thomas, E.M. Warner (2019), Vol. 523 of *ASP Conf. Ser.*, p. 521, arXiv:1812.03248
- [5] M. Gower et al., *Adding Workflow Management Flexibility to LSST Pipelines Execution*, in *Astronomical Data Analysis Software and Systems XXXII* (2023), Vol. in press of *ASP Conf. Ser.*, arXiv:2211.15795
- [6] W. O'Mullane, *RTN-001 - Data Preview 0: Definition and planning*. (2021), Vera C. Rubin Observatory Technical Note, <https://rtn-001.lsst.io/>
- [7] G. Mainetti, F. Hernandez, F. Jammes, Q. Le Boulc'h, *Experience deploying an analysis facility for the Rubin Observatory's Legacy Survey of Space and Time (LSST) data*, in *Proc. of CHEP 2023* (to appear)
- [8] W. O'Mullane, Y. Alsayyad, H.F. Chiang, F. Economou, M. Graham, L. Guy, H. Lin, F. Mueller, T. Jenness, C. Slater et al., *RTN-041 - Data Preview 0.2 and Operations rehearsal for DRP*. (2023), Vera C. Rubin Observatory Technical Note, <https://rtn-041.lsst.io/>
- [9] LSST Dark Energy Science Collaboration et al., arXiv e-prints (2021), arXiv:2101.04855
- [10] Y. Babuji, A. Woodard, Z. Li, D.S. Katz, B. Clifford, R. Kumar, L. Lacinski, R. Chard, J. Wozniak, I. Foster et al., *Parisl: Pervasive Parallel Programming in Python*, in *28th ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC)* (2019), <https://doi.org/10.1145/3307681.3325400>
- [11] *Slurm Workload Manager*, <https://slurm.schedmd.com>
- [12] [https://github.com/LSSTDESC/gen3\\_workflow](https://github.com/LSSTDESC/gen3_workflow)
- [13] <https://github.com/lsst-dm/dp02-processing/>
- [14] F. Hernandez, Q. Le Boulc'h, J. Bosch, H.F. Chiang, J. Chiang, T. Jenness, B. White, *RTN-029 - Procedure for creating a butler repository at FrDF for Data Preview 0.2* (2022), Vera C. Rubin Observatory Technical Note, <https://rtn-029.lsst.io/>
- [15] *Apptainer*, <https://apptainer.org>
- [16] *CernVM-FS*, <https://cernvm.cern.ch/fs>
- [17] *CephFS*, <https://docs.ceph.com/en/latest/cephfs/>
- [18] *PostgreSQL relational database*, <https://www.postgresql.org>
- [19] T. Jenness, *DMTN-177 - Limiting Registry Access During Workflow Execution* (2023), Vera C. Rubin Observatory Data Management Technical Note, <https://dmtn-177.lsst.io/>
- [20] *SQLite*, <https://www.sqlite.org>
- [21] T. Mkrtchyan, O. Adeyemi, P. Fuhrmann, V. Garonne, D. Litvintsev, A. Millar, A. Rossi, M. Sahakyan, J. Starek, S. Yasar, *dCache - storage for advanced scientific use cases and beyond*, in *EPJ Web of Conferences* (2019), Vol. 214, p. 04042
- [22] E. Karavakis, W. Guan, Z. Yang, T. Maeno, W. Torre, J. Adelman-McCarthy, F. Barreiro Megino, K. De, R. Dubois, M. Gower et al., *Integrating the PanDA Workload Management System with the Vera C. Rubin Observatory*, in *Proc. of CHEP 2023* (to appear)