Draft version February 25, 2025
Typeset using LATEX **twocolumn** style in AASTeX631

Rubin Observatory Data Facilities

Adam Bolton,[1] William O'Mullane,[2] Kian-Tat Lim,[1] Fabio Hernandez,[3] Mark G. Beckett,[4] Michelle Butler,[5]
Tim Jenness,[6] Orion Eiger,[7,*] Peter Love,[8] Timothy Noble,[9] Stephen Pietrowicz,[5]

Rubin Observatory Data Management Team

[1] *SLAC National Accelerator Laboratory, 2575 Sand Hill Rd., Menlo Park, CA 94025, USA*
[2] *Vera C. Rubin Observatory, Avenida Juan Cisternas #1500, La Serena, Chile*
[3] *CNRS, CC-IN2P3, 21 avenue Pierre de Coubertin, CS70202, F-69627 Villeurbanne cedex, France*
[4] *Institute for Astronomy, University of Edinburgh, Royal Observatory, Blackford Hill, Edinburgh EH9 3HJ, UK*
[5] *NCSA, University of Illinois at Urbana-Champaign, 1205 W. Clark St., Urbana, IL 61801, USA*
[6] *Vera C. Rubin Observatory Project Office, 950 N. Cherry Ave., Tucson, AZ 85719, USA*
[7] *Kavli Institute for Particle Astrophysics and Cosmology, SLAC National Accelerator Laboratory, Stanford University, Stanford, CA 94025, USA*
[8] *Lancaster University, Lancaster, UK*
[9] *Science and Technology Facilities Council, Rutherford Appleton Laboratory, Harwell, UK*

(Dated: February 25, 2025)

ABSTRACT

As the Commissioning Execution Plan (LSE-390) says, ”The project team shall deliver all reports documenting the as-built hardware and software including: drawings, source code, modifications, compliance exceptions, and recommendations for improvement.” As a first step towards the delivery of documents that will describe the system at the end of construction, we are assembling teams for producing of the order 40 papers that eventually will be submitted to relevant professional journals. The immediate goal is to accomplish all the writing that can be done without data analysis before the data taking begins, and the team becomes much more busy and stressed.

This document provides the template for these papers.

1. INTRODUCTION

The Data Facilities support Rubin Data Management by providing the underlying infrastructure used for productions and...

(for the layperson, what are we trying to do) (data rights, security requirements)

2. DATA STORAGE

2.1. *Identity Management and Security*

2.2. *Archives*

2.2.1. *Images*

2.2.2. *Engineering/Facilities Database (EFD)*

2.2.3. *Large File Annex (LFA)*

2.2.4. *Summit Database*

2.2.5. *Alert Production Products*

2.2.6. *Prompt Products Database*

2.2.7. *Data Release Products (includes ConsDB snapshot)*

2.3. *Storage Types*

2.3.1. *POSIX*

2.3.2. *S3*

2.3.3. *WebDAV*

2.3.4. *Tape*

2.4. *Databases*

2.4.1. *SQL*

2.4.2. *InfluxDB*

3. DATA MOVEMENT AND MANAGEMENT

3.1. *Long-Haul Networks*

3.1.1. *Summit to USDF*

3.1.2. *USDF to Europe*

3.2. *Rucio and Butler Integration*

3.3. *s3daemon*

[*] SLAC

3.4. *Disaster Recovery*

4. MULTI-SITE PROCESSING

4.1. *CVMFS Software Distribution*

4.2. *Embargo and Unembargo*

4.3. *Processing Middleware: The Butler and BPS*

4.4. *Campaign Management*

The term “Campaign” is defined in DMTN-181 as “a data processing activity that accomplishes a particular goal” (Lim et. al, in prep.). As such, Campaign Management refers to the process of, and the team which manages, such data processing efforts. Each data release consists of a large-scale campaign, though much shorter Data Release Processing and Alert Processing campaigns are regularly performed for development, testing, and quality assurance purposes. The Science Pipelines team has made extensive use of reprocessing campaigns of the HSC-RC2 and DC2 test-med-1 datasets in particular (“How DM uses precursor datasets and stories from HSC[sic]”, Yusra AlSayyad, Rubin Project Community Workshop, 10 August 2023). The larger the campaign is, the more difficult the task of managing it. While small campaigns can be managed easily with batch jobs, larger campaigns require data to be divided up into consistent groups, distributed to different data facilities depending on presence of data products and availability of compute resources, transferred to one location for global calibration steps, transferred back out for more parallelizable processing, and generally divided and re-combined over a number of different axes and processing steps. These groups also must be tracked in terms of what happened during the processing; did it halt due to errors (and if so, what kind)? Did it finish successfully, and is it ready to be transferred, move on to the next step, or perhaps waiting for additional data to arrive?

Campaigns rely on multiple layers of infrastructure: nearly every layer described in this paper is relevant to the execution of a campaign, though some are more immediate than others. Data products must be read from and ingested to a Butler; BPS is used to launch and manage running jobs, along with a workflow management system such as PanDA (cite) or HTCondor (also cite). Quantum Graphs are the DAGs which actually represent Science Pipeline execution (citation), and can be used after-the-fact to track execution outcomes (“Using the Processing Summaries from the Quantum Provenance Graph”, Orion Eiger, Science Pipelines Team Meeting, 9 October 2024). The Science Pipelines algorithms are another layer. Problems can arise at any level, and with so many layers, troubleshooting can be incredibly complex.

Campaigns must be run end-to-end, and advance as quickly as possible in order to keep up with real-time data volumes of the type that LSSTCam will provide. This is a difficult prospect in the scheme of an entire Data Release Production, and cannot be carried out by a person with batch scripts alone.

Campaign management is then, in essence, the orchestration of all of this. Because the process is so complex, we have developed tooling called cm-service (how to cite a github repository?). cm-service defines a campaign as a “graph of graphs” (“Status of the Campaign Management Tools,” Toby Jennings and Orion Eiger, Lyon Data Facilities Summit 2025) that includes Pipeline stages and other functions required for managing the processing of many groups of data. cm-service runs in a Phalanx-deployed Kubernetes environment where it can access a database and one or more workflow management system, and provides an always-on loop which inspects active campaigns in order to update and evolve their states. For data release productions, cm-service will manage splitting data into groups based on sharding dimensions, collecting and pausing campaigns at relevant points for data calibration and validation purposes, streamlining the process of troubleshooting and reporting, and setting nodes of campaign processing to appropriate status. cm-service will need to be aware of data transport and locality so that processing can be orchestrated at the right place at the right time. While these are difficult requirements, cm-service is ready to rise to the occasion and will be adopted pending design review (hopefully this line will be “was adopted and is in use as of [date]” at date of publication).

(implications of facility setup on CM)

Campaigns are carried out by two individuals, a *pilot* and a *co-pilot* who hand off responsibilities to each other fluidly. Pilots are supported by the Campaign Management team for troubleshooting, reporting and monitoring, as well as the respective teams for layers of infrastructure and pipeline software as they become relevant. Because carrying out a campaign is such an involved process, detailed monitoring efforts are also necessary to inform decisions in real-time.

5. MONITORING

5.1. *Compute*

(utilization, memory, etc.)

5.2. *Networking*

(bandwidth, latency, errors/drops, etc.)

5.3. *Storage*

(size, bandwidth, quotas, etc.)

5.4. *Applications and Service Logs*

6. DEVELOPER RESOURCES

6.1. *USDF RSP*

6.2. *Interactive Nodes*

6.3. *Build and Test System*

7. DATA ACCESS CENTER(S)

7.1. *US DAC and Hybrid Model*

7.2. *Chilean DAC*

7.3. *IDACs*

8. MANAGEMENT AND COORDINATION FOR OPERATIONS

8.1. *Planning*

(sizing models, etc.)

8.2. *Communications Mechanisms*

8.3. *Incident Response*

9. FACILITY SPECIFICS

9.1. *Summit Data Facility*

(instantiation of the rest here)

9.1.1. *Commissioning Cluster*

9.1.2. *Teststands*

9.2. *USDF*

9.3. *French Data Facility*

9.4. *UK Data Facility*

9.4.1. *Lancaster*

9.4.2. *Rutherford Appleton Lab*

REFERENCES