



Vera C. Rubin Observatory
Rubin Observatory Operations

Bulk Data Transfer Policies and Procedures

Adam Bolton, Others TBD

RTN-086

Latest Revision: 2024-10-31

DRAFT



Abstract

This document describes policies and procedures for transfer of large volumes of Rubin-LSST scientific data to other sites. It is focused primarily on transferring science data products to Independent Data Access Centers and other analysis facilities. Transfer between Rubin Data Facilities in the course of multi-site data-release production is out of scope for this document.

Draft

Change Record

Version	Date	Description	Owner name
1	YYYY-MM-DD	Unreleased initial draft for discussion and comment	Adam Bolton

Document source location: <https://github.com/lstt/rtn-086>

Draft

Contents

1	Context and motivation	1
2	Anticipated use cases	1
3	Proposed technical solution	3
4	Policy considerations	5
5	Commissioning bulk-transfer capabilities via Rubin Data Previews	6
A	References	7
B	Acronyms	8

Bulk Data Transfer Policies and Procedures

1 Context and motivation

Rubin Observatory's Legacy Survey of Space and Time (LSST) will generate an unprecedented multi-petabyte dataset of astronomical images, catalogs, and associated data products. A detailed description of these data products is given in LSE-163, an itemized summary of their volume over time is given in RTN-003, and a detailed sizing model is found in DMTN-135.

The primary repository for Rubin data products is the US Data Facility (USDF), hosted at SLAC National Accelerator Laboratory. The primary user interface for the scientific community to access Rubin-LSST data is the Rubin Science Platform (RSP). The RSP vision is described in LSE-319, and detailed RSP requirements are collected in LDM-554. The RSP is hosted in the commercial cloud and is supported by database query and data-access services at the USDF. Together, the cloud-based RSP and the USDF data backend constitute the US Data Access Center (US DAC).

While the US DAC will cover a large fraction of LSST science-analysis use cases under the model of "bring your analysis to the data", some use cases may require the export of LSST data to other locations, hence the need for a bulk download capability.

This document describes current plans to enable bulk download and transfer of LSST data products. It is expected to evolve based on ongoing discussions and lessons learned through Rubin Data Previews.

2 Anticipated use cases

This section provides a non-exhaustive list of anticipated use cases for bulk download. What qualifies as "bulk" is not defined here quantitatively. Qualitatively we take it to refer to the transfer of a relatively large volume of LSST data products to another facility for analysis or reprocessing.

- **Transfer to IDACs under the in-kind program:** Rubin-LSST includes a program of in-kind contributions from international participant communities in exchange for data

rights for an enumerated set of PIs. Several of these contributions take the form of Independent Data Access Centers (IDACs), which will provide the entire LSST science community with access to mirrored copies of LSST data products along with various additional local computational and user-support resources. The in-kind program may expand in the future to encompass IDAC contributions from US domestic groups and organizations, or from additional international partners. IDACs are expected to support certain benchmark capabilities that are described in RTN-003. Facilities supporting the subsequent use cases outlined immediately below are not assumed to serve as IDACs according to the criteria of RTN-003, although in principle they are not precluded from doing so.

- **Coordinated analysis by LSST Science Collaborations at other facilities:** We anticipate that the LSST Science Collaborations will obtain significant allocations of computing resources at other facilities, and will need to stage copies of LSST data products at those locations. We currently expect that the Dark Energy Science Collaboration (DESC) will conduct many of its analysis campaigns at NERSC, and that a full copy of the latest LSST data releases will be copied to that location. Other science collaborations may obtain allocations at other facilities, also requiring bulk data transfer.
- **Community mirroring of LSST data at other analysis facilities:** NSF-funded supercomputing facilities such as the Texas Advanced Computing Center (TACC) and DOE-funded centers such as the Argonne and Oakridge Leadership Computing Facilities (ALCF and OLCF) provide access to proposal-driven allocations of computing resources that far exceed what Rubin can provide to the community. To enable the community to leverage these resources for the analysis of Rubin-LSST data, mirroring of high-value LSST data products to these facilities may be of interest. (This use case may overlap with the Science Collaboration analysis use case above.)
- **Campus-level download:** Some colleges and universities may plan to devote significant campus-level computing resources to hosting and analyzing LSST data, and may request bulk download to their campus data centers.
- **Individual power users:** Individual Rubin science users may wish to deploy computing resources at their individual disposal, potentially requesting download of significant quantities of Rubin data.

For the purposes of this document, “bulk download” refers specifically to outbound transfers to facilities outside the scope of Rubin operations. Data transfers to the French and UK Data

Facilities (FrDF and UKDF) are not called out above because they are integrated into the Rubin data-processing model. The transfer of data products to the Chilean DAC is likewise not included because that process is also internal to Rubin operations. Finally, this document does not consider the possibility of *upload* of derived data products into the Rubin Data Facilities.

3 Proposed technical solution

Rubin employs a distributed computing model, with significant shares of each annual data release being processed at each of the USDF, FrDF, and UKDF. The multi-site processing model is described in DMTN-213. This approach requires a coordinated distribution of input and output data across the three sites, for which Rubin has adopted the Rucio distributed data-management software system (Barisits et al., 2019)¹. The baseline plan is to use Rucio also to enable bulk download of science data products.

In addition to the benefit of aligning bulk download tools with distributed processing tools, there are several other reasons supporting the choice of Rucio for bulk download of LSST data:

1. It is open-source and free to use.
2. It has been developed to support large-scale high-energy physics experiments with which Rubin shares many data-management similarities.
3. Many Rubin personnel and collaborators have significant experience with Rucio.
4. Transfers can be managed from a central location according to a flexible set of rules.
5. Rucio includes capabilities to monitor the success of data transfers and the integrity of the data at remote sites.
6. Transfers can be orchestrated not only *to* other facilities, but also *between* other facilities, potentially sharing the load and reducing the impact of network bandwidth limits out of the USDF. This creates the possibility of a tiered approach with a smaller number of IDACs serving as regional redistribution sites serving a larger network of downstream IDACs.

¹<https://rucio.cern.ch>

The current plan is that bulk download and transfer of Rubin data will be orchestrated through a Rucio service hosted at the USDF, and that receiving sites will be configured as *Rucio Storage Elements* (RSEs) in a data distribution network employing Rucio's *data identifier* (DID), *subscription*, and *replication rule* features.

Addition elements of the baseline plan are:

- Data transfers will be orchestrated in conjunction with annual data releases.
- Transfer sites will need to provide their own local expertise in configuring their local RSE.
- Rubin staff will convene meetings and communication channels to provide support to transfer sites, and to facilitate transfer sites' opportunities to learn from one another's experiences.
- Rules for distribution to RSEs will be implemented and maintained centrally at the USDF.
- To the greatest extent possible, DIDs will be configured to tag pre-defined collections of data, based for example on data release association, data-product type, and standardized definitions of the contents of a "Lite IDAC" or "Full IDAC". The current expectation is that these collections will be limited to subsets of the files in a given data release, and will not support distributing subsets of the contents of individual files.
- Transfer sites may install and configure local copies of Rubin middleware services (Data Butler, DMTN-229) and RSP but are not required to do so (unless required under the terms of an in-kind agreement).

There are several outstanding technical challenges that may affect the practical utility of Rucio to support Rubin bulk downloads:

1. Sites that wish to deploy Rubin middleware and science-platform services in conjunction with their copies of Rubin data will need to integrate Rucio with their local Rubin Data Butler databases. Tools for this integration are currently under development.
2. Rubin data releases are characterized by a relatively large number of relatively small files, compared to the high-energy experimental datasets to which Rucio was originally applied. The impact of this on Rucio data-transfer performance is currently being investigated, along with potential mitigations.

3. Normal operation of a Rucio Storage Element requires the use of certificates and authorization mechanisms that may be inconsistent with local site policies. Alternatives and workarounds are available and will be further explored in the context of Rubin Data Previews.

The initial focus of Rubin bulk download process development will be to support coordinated transfers of data to institutions that will in turn serve large constituencies of users. It may eventually be possible to support smaller scale user-driven “pulls” of data. Several options are under consideration for this, subject to technical feasibility and available support resources:

- Use of the Rucio Command Line Interface (CLI) to enable Rucio-based pulls of data.
- Download capability through Rucio-Globus integration planned to be implemented by an ongoing project through the Argonne High-Energy Physics Center for Computational Excellence (HEP-CCE)².
- Download based purely on Globus and/or HTTPS protocol

Any of the above scenarios will require authentication, and users or institutions accessing the data through these channels will be bound by Rubin data rights restrictions.

4 Policy considerations

Rubin Observatory's data rights policy is detailed in RDO-013. All bulk download of LSST data is subject to this policy. Individuals, institutions, and communities initiating or receiving bulk transfers of LSST data are responsible for familiarizing themselves with this policy and ensuring that their use of data delivered through bulk transfer channels complies with it.

Bulk downloads of Rubin data will only be arranged for institutions whose user community includes Rubin data-rights holders, and who can present a credible plan for (a) restricting access to users with Rubin data rights and (b) not re-exporting data to locations where it can be accessed by individuals without data rights. Institutions may wish to make use of Rubin data-rights authentication-as-a-service through the US DAC as described in DMTN-253.

²<https://www.anl.gov/hep-cce>

To maximize the timely availability of Rubin data releases to the scientific community, bulk transfers of the data products for a given data release may be made in advance of the actual release date. In these cases, bulk-transfer recipient sites must embargo the data from scientific use until the official release date.

The availability of Rubin operations staff to support bulk download is limited. Network bandwidth is likewise limited (see RTN-003 for estimates of network resource demands.) Effort to set up and support bulk download capabilities will be prioritized into three tiers:

- **Tier 1:** Transfer to IDACs under the in-kind program and to other facilities supporting major Science Collaboration analyses.
- **Tier 2:** Transfer to facilities that can provide open-access computing resources (for example through a peer-reviewed allocation process) in conjunction with hosted copies of LSST data. Facilities in Tier 2 that are willing and able to serve as intermediary RSEs for onward distribution to other mirror sites may be given additional priority.
- **Tier 3:** All other facilities and individuals.

IDACs and other institutions receiving bulk transfers will be encouraged to share knowledge and expertise with one another directly, to maximize the collective capacity of the Rubin stakeholder community to support bulk data transfer.

5 Commissioning bulk-transfer capabilities via Rubin Data Previews

This document describes a baseline plan for supporting bulk download of Rubin data. Establishing a detailed procedure will require collaboration between Rubin Data Facilities staff and IDACs and other mirror sites throughout the Data Previews of the Rubin Early Science Program described in RTN-011. Data Facilities staff will solicit expressions of interest and initiate coordination meetings during the DP1 pre-release phase. Staff availability will be limited by the need to focus on the priority of commissioning Rubin’s multi-site processing framework. Interested parties should join the #idac-coordination-group channel in the LSSTC Slack space, which is being used as a central clearinghouse for conversation and planning.

A References

- [DMTN-253]**, Allbery, R., 2024, Science Platform authentication for IDAC users, URL <https://dmtn-253.lsst.io/>,
Vera C. Rubin Observatory Data Management Technical Note DMTN-253
- Barisits, M., Beermann, T., Berghaus, F., et al., 2019, Computing and Software for Big Science, 3, 11, URL <https://doi.org/10.1007/s41781-019-0026-3>
- [RDO-013]**, Blum, R., the Rubin Operations Team, 2020, Vera C. Rubin Observatory Data Policy, URL <https://lsst/RDO-013>,
Vera C. Rubin Observatory RDO-013
- [LDM-554]**, Dubois-Felsmann, G., Ciardi, D., Mueller, F., Economou, F., 2019, Data Management LSST Science Platform Requirements, URL <https://ldm-554.lsst.io/>,
Vera C. Rubin Observatory Data Management Controlled Document LDM-554
- [RTN-011]**, Guy, L.P., Bechtol, K., Bellm, E., et al., 2024, Rubin Observatory Plans for an Early Science Program, URL <https://rtn-011.lsst.io/>,
Vera C. Rubin Observatory Technical Note RTN-011
- [DMTN-229]**, Jenness, T., 2022, The Vera C. Rubin Observatory Data Butler and Pipeline Execution System, URL <https://dmtn-229.lsst.io/>,
Vera C. Rubin Observatory Data Management Technical Note DMTN-229
- [LSE-319]**, Jurić, M., Ciardi, D., Dubois-Felsmann, G., Guy, L., 2019, LSST Science Platform Vision Document, URL <https://lse-319.lsst.io/>,
Vera C. Rubin Observatory LSE-319
- [LSE-163]**, Jurić, M., Axelrod, T., Becker, A., et al., 2023, Data Products Definition Document, URL <https://lse-163.lsst.io/>,
Vera C. Rubin Observatory LSE-163
- [DMTN-213]**, Lim, K.T., 2023, Multi-Site Data Release Processing Using PanDA and Rucio, URL <https://dmtn-213.lsst.io/>,
Vera C. Rubin Observatory Data Management Technical Note DMTN-213

[RTN-003], O'Mullane, W., Willman, B., Graham, M., et al., 2021, Guidelines for Rubin Independent Data Access Centers, URL <https://rtn-003.lsst.io/>, Vera C. Rubin Observatory Technical Note RTN-003

[DMTN-135], O'Mullane, W., Dubois, R., Butler, M., Lim, K.T., 2023, DM sizing model and cost plan for construction and operations., URL <https://dmtn-135.lsst.io/>, Vera C. Rubin Observatory Data Management Technical Note DMTN-135

B Acronyms

Acronym	Description
CLI	Command Line Interface
DAC	Data Access Center
DESC	Dark Energy Science Collaboration
DMTN	DM Technical Note
DOE	Department of Energy
DP1	Data Preview 1
FrDF	French Data Facility
HEP	High Energy Physics
IDAC	Independent Data Access Center
LDM	LSST Data Management (Document Handle)
LSE	LSST Systems Engineering (Document Handle)
LSST	Legacy Survey of Space and Time (formerly Large Synoptic Survey Telescope)
LSSTC	LSST Corporation
NERSC	National Energy Research Scientific Computing Center
NSF	National Science Foundation
OPS	Operations
RDO	Rubin Directors Office
RSE	Rucio Storage Element
RSP	Rubin Science Platform
RTN	Rubin Technical Note
SLAC	SLAC National Accelerator Laboratory
TACC	Texas Advanced Computing Center
TBD	To Be Defined (Determined)

UK	United Kingdom
UKDF	United Kingdom Data Facility
US	United States
USDF	United States Data Facility

Draft