

The Vera C. Rubin Observatory Data Preview 2

VERA C. RUBIN OBSERVATORY TEAM,¹ ERIC C. BELL¹,² PEDRO H. BERNARDINELLI¹,³ AUDREY BUDLONG¹,⁴
RODOLFO CANESTRARI¹,⁵ NEVEN CAPLAR¹,² COLIN ORION CHANDLER¹,^{6,2,7} JOHANN COHEN-TANUGI¹,⁸
CYRILLE DOUX¹,⁹ SIEGFRIED EGGL¹,¹⁰ ANTHONY ENGLERT¹,¹¹ SHENMING FU¹,¹² MELISSA L. GRAHAM¹,^{2,3}
LEANNE P. GUY¹,¹³ MATTHEW J. HOLMAN¹,¹⁴ MINHEE HYUN¹,¹⁵ MARIO JURIC¹,³ YIJUNG KANG¹,^{12,13}
ARUN KANNAWADI¹,^{16,17} KSHITIJ KELKAR¹,¹³ PIERRE-FRANÇOIS LÉGET¹,¹⁷ SHUANG LIANG,¹⁸ KIAN-TAT LIM¹,¹⁸
SEAN PATRICK MACBRIDE¹,¹⁹ SAMUEL MESQUITA¹,⁹ JAMES R. MULLANEY¹,²⁰ KATE NAPIER¹,¹²
WILLIAM O’MULLANE¹,¹³ AASHAY PAI¹,²¹ KARLA PEÑA RAMÍREZ¹,¹³ ANDRÉS A. PLAZAS MALAGÓN¹,^{18,12}
GABRIELE RODEGHIERO¹,²² LUCA ROSIGNOLI¹,^{23,22} IAN S. SULLIVAN¹,² ALESSIO TARANTO¹,^{22,23} AND
DAN S. TARANU¹,¹⁷

¹NSF-DOE Vera C. Rubin Observatory / NSF NOIRLab, 950 N. Cherry Ave., Tucson, AZ 85719, USA

²University of Washington, Dept. of Astronomy, Box 351580, Seattle, WA 98195, USA

³Institute for Data-intensive Research in Astrophysics and Cosmology, University of Washington, 3910 15th Avenue NE, Seattle, WA 98195, USA

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

⁵INAF Istituto di Astrofisica Spaziale e Fisica Cosmica di Palermo, Via Ugo la Malfa 153, 90146, Palermo, Italy

⁶LSST Interdisciplinary Network for Collaboration and Computing, Tucson, USA

⁷Department of Astronomy and Planetary Science, Northern Arizona University, P.O. Box 6010, Flagstaff, AZ 86011, USA

⁸LPCA, Université Clermont-Auvergne, CNRS/IN2P3, Clermont-Ferrand, France

⁹Université Grenoble Alpes, CNRS/IN2P3, LPSC, 53 avenue des Martyrs, F-38026 Grenoble, France

¹⁰University of Illinois, Department of Astronomy, 1110 W. Green St., Urbana, IL 61801, USA

¹¹Department of Physics, Brown University, 182 Hope Street, Providence, RI 02912, USA

¹²Kavli Institute for Particle Astrophysics and Cosmology, SLAC National Accelerator Laboratory, 2575 Sand Hill Rd., Menlo Park, CA 94025, USA

¹³NSF-DOE Vera C. Rubin Observatory / NSF NOIRLab, Casilla 603, La Serena, Chile

¹⁴Center for Astrophysics, Harvard & Smithsonian, 60 Garden Street, Cambridge, MA 02138

¹⁵Stanford University, 450 Jane Stanford Way, Stanford, CA 94305, USA

¹⁶Department of Physics, Duke University, Durham, NC 27708, USA

¹⁷Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA

¹⁸SLAC National Accelerator Laboratory, 2575 Sand Hill Rd., Menlo Park, CA 94025, USA

¹⁹Physik-Institut, University of Zurich, Winterthurerstrasse 190, 8057 Zurich, Switzerland

²⁰Astrophysics Research Cluster, School of Mathematical and Physical Sciences, University of Sheffield, Sheffield, S3 7RH, United Kingdom

²¹Department of Astronomy and Astrophysics, University of Chicago, 5640 South Ellis Avenue, Chicago, IL 60637, USA

²²INAF Osservatorio di Astrofisica e Scienza dello Spazio Bologna, Via P. Gobetti 93/3, 40129, Bologna, Italy

²³Department of Physics and Astronomy (DIFA), University of Bologna, Via P. Gobetti 93/2, 40129, Bologna, Italy

(Dated: July 30, 2026)

ABSTRACT

We present an early version of Rubin Data Preview 2 (DP2), the second data preview from the National Science Foundation–Department of Energy Vera C. Rubin Observatory. [Early Data Preview 2 \(EDP2\)](#) comprises coadds, detection catalogs, and ancillary data products; eventually, the full DP2 will also include single-epoch images and difference images, but any analysis begun on [EDP2](#) will require no modification. DP2 is derived from observations acquired by the LSST Science Camera (LSSTCam) on the Simonyi Survey Telescope at the Summit Facility on Cerro Pachón, Chile, primarily during the on-sky commissioning campaign between 2025-04-25 and 2025-09-21, supplemented by observations taken between 2025-10-25 and 2026-01-06 that overlap the commissioning footprint. The DP2 footprint

comprises the Science Validation wide-area survey, five Deep Drilling Fields, and a number of targeted small-field regions, including Trifid and Lagoon, Prawn, M49, and New Horizons, all observed as part of the Rubin First Look campaign. Each field was imaged in up to six broad photometric bands, *ugrizy*, and coadded to produce deep imaging covering an estimated 3,000 deg². The addition of single-visit-only areas expands the total DP2 footprint to an estimated 15,000 deg², with coverage in at least one filter. The median per-visit PSF FWHM across the wide-area survey ranges from 1.17 arcsec in the *z* band to 1.26 arcsec in *g* and *r* bands, with the sharpest fields reaching $\sim$ nan". COSMOS, the deepest field, reaches estimated coadded 5σ depths of *u*=26 mag, *g*=26.8 mag, *r*=26.3 mag, *i*=26.1 mag, *z*=25.3 mag, *y*=23.9 mag. DP2 contains 790 million distinct astrophysical objects, of which X are extended in at least one band in coadds, and X solar system objects, of which X are new discoveries. DP2 is approximately 59.15 TB in size and is available to Vera C. Rubin Observatory data rights holders via the Rubin Science Platform, a cloud-based environment for the analysis of petascale astronomical data. Based on a roughly five-month primary observing baseline and covering only part of the eventual LSST footprint, DP2's area, depth, and multiband coverage nonetheless support a broad range of early science investigations ahead of LSST Data Release 1.

Keywords: Rubin Observatory - LSST

1. INTRODUCTION

The National Science Foundation (NSF)–Department of Energy (DOE) Vera C. Rubin Observatory is a ground-based, wide-field optical-near-infrared facility on Cerro Pachón in northern Chile, built to carry out the Legacy Survey of Space and Time (LSST) (Ivezić et al. 2019). The observatory is named in honor of Vera C. Rubin, whose observations in the latter half of the 20th century established the first compelling empirical case for dark matter (Rubin & Ford 1970; Rubin et al. 1980). Its 10-yr survey will image the entire visible southern sky every 3–4 nights, reaching a coadded depth of $\sim$ 27 mag in the *r*-band across an 18000 deg² footprint (Bianco et al. 2022).

At the heart of the facility is the 3.2 gigapixels-gigapixel LSST Science Camera (LSSTCam), mounted on the 8.4 m (6.5 m effective aperture) Simonyi Survey Telescope (Stalder et al. 2024). Together they deliver an instantaneous field of view of 9.6 deg² in six broadband filters, *ugrizy* (320–1050 nm), with seeing-limited image quality. The system's étendue of 319 deg² m² exceeds that of any previous optical facility by over an order of magnitude, enabling the rapid, deep, all-sky cadence that underpins Rubin's four core science themes: the nature of dark energy and dark matter; a census of the solar system; the transient and variable optical sky; and the structure of the Milky Way (Ivezić et al. 2019). An automated Data Management System processes tens of terabytes of images per night to produce science-ready data products within minutes of each observation for a global community of scientists.

This paper presents Rubin's third Data Preview, Data Preview 2 (DP2) (NSF-DOE Vera C. Rubin Observatory

2025), and the first to be based entirely on LSSTCam data. It marks the conclusion of the Rubin commissioning program and the transition to full survey operations. DP2 is a reprocessing of all science-grade exposures acquired with LSSTCam during on-sky commissioning and the initial phase of early operations, prior to the formal start of the LSST on 2026-06-29. As such, it represents a natural boundary and provides a comprehensive dataset that captures the on-sky performance of the as-built system at the end of the Rubin construction project.

Two earlier Data Previews preceded DP2 as part of the Rubin Early Science Program (Guy et al. 2026), established to provide the community with early access to data products and services needed to produce high-impact science during the period spanning commissioning through the first Data Release. Data Preview 0 (DP0) was based on simulated LSST-like data (LSST Dark Energy Science Collaboration (LSST DESC) et al. 2021) and served primarily to stand up early versions of the Rubin Science Platform (RSP) and prepare the community for data access at scale. Data Preview 1 (DP1) (NSF-DOE Vera C. Rubin Observatory 2025) delivered the first data products derived from on-sky observations using science-grade exposures acquired with Rubin Commissioning Camera (LSSTComCam) (Stalder et al. 2022, 2020; Howard et al. 2018; SLAC National Accelerator Laboratory & NSF-DOE Vera C. Rubin Observatory 2024); a 144-megapixel, single-raft instrument used during early commissioning prior to the installation of LSSTCam on the Simonyi Survey Telescope. DP2 supersedes both in scope, fidelity, and scientific depth.

DP2 contains 29756 science-grade exposures acquired over 158 nights between 2025-04-16 and 2025-09-21, cov-

ering a total sky area of 15,000 deg² across 85 fields. The data products include raw and calibrated single-epoch images, coadded images, difference images, source detection catalogs, and a range of derived data products. The release totals approximately 59.15 TB and contains around 790 million distinct astronomical objects. Full release documentation is available at <https://dp2.lsst.io>.

Because DP2 is drawn from a commissioning and early-operations data rather than from the more uniform and systematic LSST, its images and derived products reflect a system that was still being commissioned and tuned. LSSTCam’s greater scale and complexity relative to LSSTComCam introduced systematics with no direct counterpart in DP1. For much of the campaign the Active Optics System (AOS) closed loop converged on a slightly mis-calibrated focus target, imprinting a spatially varying defocus across the focal plane that was only corrected on sky in December 2025 and that remains visible in the point-spread function (PSF) of a large fraction of DP2 exposures. Delivered image quality and PSF ellipticity did not consistently meet LSST design requirements in wide-area survey-mode observations by the close of commissioning; and stray- and scattered-light artifacts, sky-background systematics in crowded fields, and a chromatic PSF dependence²⁴ were identified and only partially mitigated in DP2 processing. These effects, their diagnosis, and the mitigations applied are described in §2 and §5.

Unlike the systematic, rolling revisit pattern that will characterize the LSST, DP2’s observing cadence is opportunistic: fields and filters were prioritized to serve AOS and Science Pipelines commissioning rather than to build a representative time series. The Science Validation survey footprint was repeatedly re-scoped over the course of the science validation campaign due to the competing demands of engineering tests and filter availability, which was constrained by hardware issues for part of the campaign. Visit counts and time sampling therefore vary substantially and idiosyncratically from field to field, and DP2 should not be read as a preview of LSST’s eventual survey cadence, revisit statistics, or depth-accumulation pattern.

DP2 is being delivered in two phases across the Jul–Dec 2026 window. The first phase, Early DP2 (EDP2), comprises the deep coadd images and all associated catalogs, including new cell-based coadds and the first releases of the Object Shear (§3.2.3) shape-measurement catalog and the Isolated Star Stellar Motions catalog (§3.2.2). EDP2 is targeted for 2026-07-27. The sec-

ond phase, DP2, adds the visit-level single-epoch and difference images, completing the full release, and is expected in the Oct–Dec 2026 range. Data products delivered as EDP2 will be incorporated into DP2 unchanged, without additional reprocessing (Guy et al. 2026). This phased delivery follows directly from Rubin’s hybrid cloud architecture (O’Mullane et al. 2024a), in which science data products are stored at the United States Data Facility (USDF) and served to users via the US Data Access Center (DAC) in Google Cloud. Deep coadd images and catalogs are hosted directly in Google Cloud and are therefore unaffected by hybrid-cloud readiness, allowing them to ship as EDP2. Visit-level images, however, must be streamed from the USDF, and the necessary capacity is not yet fully in place, introducing timing uncertainty for the second phase. Delivery of the Prompt Data Products follows the same logic. The Prompt Products DataBase (PPDB), Rubin’s real-time catalog of transient, variable, and moving-object detections, does not depend on USDF–Google Cloud bandwidth and is therefore prioritized as the second Early-Operations delivery milestone, following EDP2 and preceding the full DP2 release, expected in the Sep–Oct 2026 range. Prompt images (raw images, Processed Visit Images, and difference images), which do require that bandwidth, have been deprioritized relative to EDP2 and the PPDB and are expected to become available only in the Oct–Dec 2026 range, following both.

DP2, like all Rubin Data Previews and Data Releases, is subject to a 2 yr proprietary period; during this time access is restricted to LSST data rights holders (Blum & the Rubin Operations Team 2020) via the RSP (Jurić et al. 2019), Rubin’s cloud-based data science platform (Jurić et al. 2019). After the proprietary period, DP2 will be made public, though access through Rubin DACs in the US and Chile will remain limited to data rights holders (Blum & the Rubin Operations Team 2020). Alternative public access mechanisms have not yet been finalized.

The structure of this paper is as follows. Section 2 describes the observatory system and its status at the close of commissioning, the fields included in DP2, and the observing strategy employed. Section 3 summarizes the data products in the release. The processing pipelines are described in Section 4, followed by data validation and performance characterization in Section 5. Section 6 describes the RSP. Section 7 outlines Rubin’s community support model. A summary and outlook toward the start of formal LSST Data Releases is given in Section 8.

All magnitudes are quoted in the AB system (Oke & Gunn 1983) unless otherwise specified.

²⁴ LSSTCam has no atmospheric dispersion corrector

2. COMMISSIONING WITH LSSTCAM

3. OVERVIEW OF THE CONTENTS OF RUBIN DP2

In this section we describe the data products that are included in DP2. Data rights holders can access all of these data products via one or more of the data services outlined in §6. For programmatic reasons, the precise names of the data product can differ slightly depending on which data service is being used to retrieve them. For this reason, we use descriptive names throughout this section and provide their equivalent names used in the data services in Table 1. We note that some data products are not available in all data services; such cases are indicated in Table 1.

As with DP1, all the data products described in this section fall under one of five categories:

- Science Images, including pre-processed Raw Images, Visit Images, and Deep Coadd images;
- Catalogs of astrophysical Sources and Objects detected and measured in the aforementioned science images. We also provide the reference catalog that was constructed to astrometrically and photometrically calibrate the data during processing;
- Maps, which provide non-science-grade visualisations of the data;
- Ancillary Data Products, such as the values of the configuration parameters that were used during data processing, processing logs, and calibration data products;
- Metadata, in the form of tables containing information about each visit and visit image, such as visit pointing coordinates, exposure times, and data quality summary statistics.

We expect that the Science Images and Catalogs will be the most heavily used categories. We expect Science Images and Catalogs to be the most commonly used categories for science, while Maps, Ancillary Data Products, and Metadata remain important for facilitating science exploitation via visual inspection, as well as for transparency, EPO, media, and citizen-science activities such as Zooniverse. While the categories of data products in DP2 are the same as in DP1, this second Data Preview contains a wider range of data product types compared to DP1. This is principally due to the science processing pipeline being in a more mature state compared to the first data preview. In the following sections we clearly indicate those data product types which are new to DP2.

As with DP1, the science processing pipelines use a skymap to structure how processing is organized spatially. The skymap divides the sky into a series of equal-area tracts, which are then subdivided into a number of patches. The tracts are arranged in rings of constant decl., starting at the southern pole and having an index number that increases first by increasing R.A. and then by increasing decl.. The skymap that was generated for DP2 processing contains a total of 18938 tracts spanning the whole sky, 2174 of which were covered during the science validation observing campaign, including the small field surveys. Each DP2 tract covers approximately 2.8 deg^2 , and contains 10×10 patches, both the same as for the DP1 skymap. Tracts are defined to have equal area, and so the amount by which neighboring tracts overlap varies with decl., with tracts near the celestial poles overlapping more than those near the celestial equator. The minimum area of overlap between tracts is $1'0$. Within a tract, the area of overlap between neighboring patches is a constant $60'0$. As a result, each patch covers an area of 0.034 deg^2 , which is larger than the tract area divided by the number of patches in a tract. Excluding the overlap region, the area covered uniquely by a single patch is 0.028 deg^2 .

Although the tract and patch configuration is unchanged between DP1 and DP2, the two skymaps differ at the cell level for Deep Coadds. Patches are further subdivided into cells for the sole purpose of enabling the construction of cell-based coadds with a continuous PSF (§3.1.3). Unlike patches, cells are not exposed to the user, and one cannot specify nor retrieve an individual “cell” from an image. Instead, cells are used during the construction of what we termed cell-based coadds (refs). A region of an input image is included in a given cell if, and only if, it covers the entire cell. This guarantees that every pixel within a cell is built from exactly the same set of input images, so the PSF is well-defined and effectively constant across the cell, with discontinuities confined to the cell boundaries. Such cell-based coaddition was not implemented during the production of DP1, and was only used to produce the Deep Coadds in DP2. Template Coadds in DP2 do not utilize cells. In the case of DP2 Deep Coadds, there are 22×22 cells per patch, including the patch overlap region (the patch overlap region is 1 cell wide). There is no overlap between neighboring cells.

3.1. Science Images

Science images are images taken with the intent of scientific analysis of astronomical objects, as opposed to calibration images (e.g., bias, darks, flats) or engineering images, which serve to calibrate or test the instrument

Table 1. Summary of dataset types released in EDP2, together with the names by which they are referred to within the QServ database and the Data Butler.

Dataset Type	QServe Name	Butler Name
Deep Coadd	–	deep_coadd
Object	dp2.Object	object
Isolated Star Stellar Motions	dp2.IsolatedStarStellarMotions	isolated_star_stellar_motions
Object Shear	dp2.ShearObject	object_shear_all
Source	dp2.Source	source
Forced Source	dp2.ForcedSource	object_forced_source
Forced Source at DIA Object	dp2.ForcedSourceOnDiaObject	dia_object_forced_source
DIA Object	dp2.DiaObject	dia_object
DIA Source	dp2.DiaSource	dia_source
Solar System Object	dp2.SSObject	ss_object
MPCORB	dp2.mpc_orbits	–
Calibration Reference Catalog	–	the_monster_20250219
Survey Property Maps	–	deepCoadd_<PROPERTY>_consolidated_map_<STATISTIC>
HiPS Maps	–	–
Skymap	–	skyMap
Standard Passbands	–	standard_passband
Object Scarlet Models	–	object_scarlet_models
Deep Coadd Input Summary	–	deep_coadd_input_summary
Visit	dp2.Visit	visit_table
Visit Detector	dp2.VisitDetector	visit_detector_table
Visit Summary	–	visit_summary

and telescope regardless of whether they happen to be pointed at the night sky. DP2 provides all images in the acquisition and processing chain, from the raw, unprocessed science and calibration images through to fully calibrated visit images, coadds, and template images. Most science analyses will start from the Deep Coadd images and the single-epoch Visit images. While it is possible for users to recreate their own Processed Visit Images (PVIs) from the raw and calibration data, we do not expect anyone will need to do so.” One important difference between DP1 and DP2 is the format of the Science Images (Bosch 2026). In the case of DP2, the format of the datasets are such that they are more compatible with external packages, in particular Astropy. It is, however, straightforward to convert back to “legacy” DP1-style datasets for backwards-compatibility.

3.1.1. Raw Images

Note: Raw Images are NOT included in EDP2, but we describe them here as it provides useful context for the dataset types that are included in EDP2. Raw images will be included in DP2; the

number of raw images in DP2 will be provided at that time.

Raw Images are those that are produced by the camera following an exposure. They are the initial input science images to the pipelines and have not undergone any processing. Each active detector in LSSTCam (§??) produces one Raw Image per exposure. Each exposure therefore results in up to 180 Raw Images, although some exposures taken during the observing campaign resulted in fewer than 180 Raw Images due to the problematic detectors described in §2 and occasional readout failures.

All LSSTCam detectors have 16 amplifier segments, arranged in two rows of eight, each reading out its own ~ 1 megapixel region in parallel to enable the fast (~ 2 second) full focal-plane readout (§??). These sections are evident in the Raw Images. Despite LSSTCam containing detectors from two different vendors, all Raw Images have the same dimensions: 4608×4096 pixels. However, the two types of detectors have slightly different numbers of overscan and pre-scan pixels, meaning the number of exposed pixels differs between detector

type. This difference is reflected in the dimensions of the Visit Images (§3.1.2), which have had the prescan and overscan regions removed during data processing.

3.1.2. Visit Images

Note: Visit Images are NOT included in EDP2, but we describe them here as it provides useful context for the dataset types that are included in EDP2. Visit Images will be included in DP2; the number of Visit Images in DP2 will be provided at that time.

Visit Images are the first scientifically useful images to be produced by the science processing pipeline. They are the result of removing instrument signatures from the raw input images, followed by multiple rounds of background subtraction and photometric and astrometric calibration (see §4). There is one Visit Image for every raw input image that successfully passed through the first stage of the science processing pipelines. As such, there are up to 180 Visit Images for each of the exposures taken during the observing campaign.

During instrument signature removal (ISR), the science processing pipeline removes the overscan regions from the input raw images. The Visit Images therefore only contain the pixels that were exposed during the observation: 4072×4000 pixels in the case of the ITL detectors, and 4096×4004 pixels in the case of the e2v detectors. Unlike raw images, Visit Images consist of three image planes: a science image plane, a variance plane, and a mask plane. The science image is the processed exposed image, whereas the variance image persists the uncertainty associated with the value of the corresponding pixel in the science image. The mask image indicates pixels that are flagged for various different reasons, such as saturated pixels, pixels associated with detected sources, bad pixels, etc. All three images contain the same number of pixels in each dimension.

3.1.3. Coadds

Note: Template Coadds are NOT included in EDP2. They will be included in DP2.

Once produced, individual Visit Images covering the same patch of sky (as defined by the skymap) in the same band are coadded to form coadd images. Two different types of coadd images are included in DP2: Deep Coadds and Template Coadds, both of which are described in the following subsections. The production of each type of coadd image follows broadly the same process, which is described in detail in §?? but, to summarize, it involves the selection of Visit Images (in terms of patch coverage and band, and image quality), the warping of those Visit Images onto a common pixel grid, and

the coaddition of those warped images. The main difference that distinguishes each type of coadd are the selection of the Visit Images. As a result, the structure of the datasets is the same for each type: each coadd contains three planes corresponding to the science image plane (i.e., the coadd), the variance plane, and the mask plane, with all three containing the same number of pixels in each dimension. Up to two different background models are included in the coadd datasets: the background that was subtracted to create the final coadd and, in the case of the *gri* band Deep Coadds, a background that can be added-back to create “Visualization Coadds” which, while not recommended for science, can be useful for visualization. Each coadd also contains a position-dependent PSF model, WCS information, plus various other metadata.

The same skymap was used for each type of coadd in DP2, meaning they all cover the same on-sky area, including some overlap with their neighbouring coadds (see the introduction to this section). A single coadd, irrespective of type, contains 3300×3300 equal-sized square pixels, including overlap; each pixel is $0''.2$ on each side. There are up to six coadds of each type per patch - one for each of the six LSSTCam bands. While the sky coverage of coadds is similar between types and bands, they are not exactly the same due to the different Visit Image selection functions for each type, and differences in sky coverages in each band.

The number of Visit Images that contribute to a coadd varies across the patch. This is true even if a patch is fully covered by all the visits that targeted that area because of the presence of chip gaps between neighbouring detectors. The Survey Property maps (see §3.3.1) can be consulted to gain an insight into the total exposure time, and thus the number of contributing Visit Images, at any point within a coadd. Since each coadd image always covers an entire patch, it is not uncommon for regions of a coadd image to have zero Visit Image coverage. This is especially the case on the edges of an observed field. Such regions of a coadd may contain seemingly valid values (i.e., not necessarily zeros or NaNs), but will instead be indicated by the NO_DATA flag in the mask image plane. It is therefore crucial that the mask image plane is consulted when analysing any type of coadd.

Deep Coadds—As the name suggests, Deep Coadds go down to the faintest magnitudes of the three types of coadd. Two basic selection criteria were used to select visits for inclusion in Deep Coadds. The first criterion is seeing-based and selects visits with a mean (over all detectors) PSF FWHM $< 1''.7$. The second criterion is depth-based and uses the concept of “effective exposure

time”, which is the exposure time that would be needed under nominal conditions to achieve the same signal-to-noise as the visit in question (Bernstein et al. 2016). Our specific selection uses the transparency component of effective exposure time, and selects visits in which the whole-system transparency was such that the transparency component of the effective exposure time was > 0.75 times the actual exposure time (corresponding to ~ 0.16 mag of extinction relative to nominal conditions). Following the selection of visits, a second round of selection is applied at the detector (i.e., Visit Image) level. These per-detector selection criteria are summarised in Table 2. Together, these selections aim to maximise the depth of the coadd while still ensuring good image quality. As highlighted at the start of this subsection, a key difference between the Deep Coadds released in DP2 versus those released in DP1 is that the DP2 Deep Coadds are cell-based whereas the DP1 Deep Coadds were not.

There are a total of 925460 Deep Coadds in DP2.

3.2. Catalogs

Catalogs contain data for detections in the images described in §3.1, except the **Calibration** catalog, which contains reference data obtained from previous surveys. Observatory-produced metadata tables are described in §3.5. Rubin adopts a naming convention in which the term “object” refers to an astrophysical object, whereas the term “source” refers to a single-epoch detection of that object. As such, sources are associated with detections in single-epoch Visit Images and Difference Images, whereas objects are associated with detections in multi-epoch coadds. Photometric quantities in catalogs are stored as fluxes rather than magnitudes, with fluxes measured in nanojanskies ($1 \text{ nJy} = 10^{-35} \text{ Wm}^{-2}\text{Hz}^{-1}$). Fluxes are preferred for multi-epoch observations, as they can be averaged across epochs, unlike magnitudes. Additionally, flux measurements on difference images (§3.1) are computed against a template, representing a flux difference. Additionally, flux measurements on difference images (§3.1) are computed relative to a template and can therefore be negative, particularly for faint sources in the presence of noise.

The catalogs described below each vary in terms of their specific columns but most contain: one or more unique identification numbers, positional information, multiple types of flux measurements (e.g., aperture fluxes, PSF fluxes, Gaussian fluxes, etc.), and a series of Boolean flags indicating characteristics such as saturation or cosmic-ray contamination for each source or object. The solar system catalogs SSObject and SSSource

deviate from this general structure in that they instead contain orbital parameters for all known asteroids.

Where applicable, quantities are prefixed with the band in which they were measured, and all measured properties are reported with their associated 1σ uncertainties. For example, `g_ra` and `g_raErr` in the Object catalog refer to R.A. and its uncertainty, measured in the *g* band. Fluxes for various apertures are provided together with an uncertainty and a flag, and named in the format `[band_]ap[size]Flux`, where `[size]` is the aperture diameter in pixels, and `[band_]` is used where applicable (e.g., in the object catalog, but not in the source catalog; see §3.2.4 and §3.2.1). For example, `g_ap03Flux`, `g_ap03FluxErr`, `g_ap03Flux_flag` provide the flux, uncertainty and a flag measured within a 3.0 pixel aperture in the *g* band. Similarly for flux measurements using difference algorithms, e.g. `g_psfFlux` provides the flux derived using the PSF model as a weight function, forced on *g* band.

All of the catalog types that were released as part of DP1 have been retained in DP2, although in some cases additional columns have been added for DP2; we highlight prominent additional columns in the subsections below. There are also a number of additional catalogs in DP2 that do not have DP1 equivalents. Throughout this section, we provide high-level overviews of the contents of each catalog, rather than listing and describing each column. Instead, full searchable catalogue schemas, including column descriptions, are provided at <https://sdm-schemas.lsst.io/>.

All catalogs except the Calibration Reference Catalog are stored in the Qserv database (§6.5.1) and accessible via TAP, with an online DP2 catalog schema available at <https://sdm-schemas.lsst.io/dp2.html>. Catalog data are also accessible via the Data Butler (see §6.2.2).

3.2.1. Object

The Object catalog contains data on all objects detected with a greater than 5σ significance in the Deep Coadds. With coadd images produced on a per-band basis, a $> 5\sigma$ detection in one or more of the bands will result in an object being included in the Object catalog. For cases where an object is detected at $> 5\sigma$ in more than one band, a crossmatching has been performed between bands to associate an object in one band with its counterpart(s) in the other bands. As such, each entry in the Object catalog associates objects across multiple bands. The objects reported in the Object catalog have also undergone deblending; in the case of blended detections, only the deblended child objects are included in the catalog. As with the Source catalog, the criterion for inclusion in the Object catalog is a $> 5\sigma$ detection in

Table 2. Per-detector PSF-quality selection cuts for inclusion in Deep Coadd images.

Parameter	Description	Numerical cut
PSF FWHM	Same FWHM calculation as the visit-level cut, evaluated per detector rather than averaged over the visit.	< 1.8 arcsec
Median ellipticity residual	Median of $\sqrt{e_1^2 + e_2^2}$ between the PSF model and measured star shapes. ^a	< 0.19
psfStarScaledDeltaSizeScatter	Robust scatter of PSF size residuals, scaled by the median PSF size.	< 0.07
psfTraceRadiusDelta	Max–min of the PSF trace radius evaluated on a grid across the unmasked detector pixels.	< 4.4 pixels
psfApFluxDelta	Max–min of the model PSF aperture flux (normalized) across the unmasked detector.	< 1.6
psfApCorrSigmaScaledDelta	Max–min of the PSF aperture correction across the detector, scaled by <code>psfSigma</code> .	< 0.15
starEMedian	Median star ellipticity, $\sqrt{e_1^2 + e_2^2}$. ^a	< 0.2 (all bands)
starUnNormalizedEMedian	Same, computed from unnormalized second moments. ^b	< 2.3 (all bands)
nPsfStar	Number of PSF-candidate stars available on the detector; below this the PSF model is not considered well-constrained.	≥ 6 per band

^a $e_1 = (I_{xx} - I_{yy}) / (I_{xx} + I_{yy})$, $e_2 = 2I_{xy} / (I_{xx} + I_{yy})$, where I_{xx} , I_{yy} , I_{xy} are the adaptive second moments of the PSF/star.

^b Unnormalized equivalent: $\sqrt{(I_{xx} - I_{yy})^2 + (2I_{xy})^2}$, i.e. without dividing by $I_{xx} + I_{yy}$.

one of the Deep Coadds prior to deblending, yet the positions and fluxes of objects are reported postdeblending. Hence, it is possible for the Object catalog to contain objects whose flux-to-error ratios are less than five in all bands.

In addition to the general information mentioned above (i.e., IDs, positions, fluxes, and flags), the **Object** catalog also includes basic shape and extendedness information. While they may be included in future Data Releases, no photometric redshifts, Petrosian magnitudes (Petrosian 1976), proper motions or periodicity information are included in the DP2 object catalogs.

The **Object** catalog contains data for 790 million objects in DP2.

3.2.2. Isolated Star Stellar Motions

The Isolated Star Stellar Motions catalog contains the proper motions and parallaxes of isolated stars. Rubin uses a specific definition for an isolated star. Candidate isolated stars must be point-like (i.e., unresolved) detections with a signal-to-noise ratio of between 8 and 1000, and must not be flagged as blended. All such candidates within a given tract are spatially matched using a 1'' matching radius across all visits and bands covering said tract. Should the average position of a candidate have no other candidate within 2'', then it is considered to be an isolated star.

The Isolated Star Stellar Motions catalog provides a best-fit position (**ra**, **dec**), proper motion in mas yr^{-1} (**raPM**, **decPM**), and parallax in mas at a reference epoch; the **raPM** parameter already includes the $\cos(\text{decl.})$ factor. The reference epoch in MJD is included as a column (**epoch**) in the catalog. An isolated star's proper motion and parallax are calculated by fitting its position across all per-visit detections, plus any detection in our calibration reference catalog (§3.2.11) that is matched to within 0'.1 of the star's average position. The catalog also contains 1σ uncertainties for each each of the five aforementioned positional parameters (five columns: **raErr**, **decErr**, **raPMErr**, **decPMErr**, **parallaxErr**), plus a further ten columns containing off-diagonal covariance parameters (***_*_Cov**, e.g. **ra_raPM_Cov**, **raPM_parallax_Cov**).

All isolated stars have an entry in the Isolated Star Stellar Motion catalog, although no attempt is made to fit the time-varying positions of stars that are detected in fewer than 3 epochs and which do not have a matched counterpart in the reference catalog. In such cases the positional information, including uncertainties, are reported as NaN. Every entry in the Isolated Star Stellar Motion catalog has a unique identification number and, in cases where a match to the reference catalog has been

identified, the unique identification of the matched reference source.

The Isolated Star Stellar Motions catalog contains 320 million entries.

3.2.3. Object Shear

The Object Shear catalog contains object shape measurements in five different shear steps or “counterfactual” images: one in which no shear (**ns**) has been applied, and four in which an artificial linear shear has been introduced to the image (**1p**, **1m**, **2p**, **2m**). The artificial shears are in four different directions: **1p** is along the image x-direction, **1m** is along the y-direction, **2p** is along the 45 degree diagonal, **2m** is along the 135 degree diagonal.

Two measurement algorithms (*gauss* and *pgauss*) are run on each of the five different counterfactuals, one for measuring shapes and the other for fluxes. The first approach is to fit a single 2D gaussian jointly across the *r*-, *i*-, and *z*-band deep coadd images, with the ellipticity components in the image x-y (**gauss_g1**) and diagonal (**gauss_g2**) directions reported in the catalog. The second approach measures 2D gaussian-weighted moments on single band (*r*, *i*, *z*) images separately, calculates the weighted sum of the raw moments across those three bands (weighted by inverse-variance), then reports the ellipticity components of that weighted sum in the image x-y (**pgauss_g1**) and diagonal (**pgauss_g2**) directions. There is one row per object detected in each counterfactual, meaning a given object could result in up to five rows, depending on how many of the counterfactuals it is detected in at a significance of $> 5\sigma$.²⁵

In addition to the gaussian shape measurements and their auto-covariances, the Object Shear table also contains object positional information, fluxes measured from the gaussian fits (*r*, *i* and *z* bands) and gaussian moments (*g*, *r*, *i* and *z* bands), PSF shapes reported in the same x-y and diagonal directions as objects, a range of different flags. Each entry in the Object Shear catalog also has a unique identification number (**shearObjectId**), that uniquely identifies a given detection in a given counterfactual. The Object Shear catalog contains 2.5 billion rows.

3.2.4. Source

The Source catalog contains the measured properties of all sources that are, prior to deblending, detected at a $> 5\sigma$ significance in each of the Visit Images (§3.1.2). The detections in the Source catalog have all undergone

deblending; in the case of blended sources, only the post-deblended sources are included in the Source catalog. It is therefore possible for the Source catalog to contain entries in which none of the different types of photometric measurements (e.g., aperture, PSF, gaussian) have flux-to-error ratios greater than five.

Each entry in the source catalog has a unique identification number (**sourceId**), positional information, various photometric measurements (aperture, PSF, gaussian), shape and extendedness information, and a series of flags. Because each entry in the source catalog is derived from a detection in a single epoch observation taken with a single filter, each entry also has an associated band.

The Source catalog contains data for 18 billion sources in DP2.

3.2.5. Forced Source

The Forced Source catalog contains forced PSF photometry measurements performed on both Difference Images (i.e., the **psfDiffFlux** column) and Visit Images (i.e., the **psfFlux** column) at the positions of all the objects in the Object catalog, to allow assessment of the time variability of the fluxes. We recommend using the **psfDiffFlux** column when generating light curves because this quantity is less sensitive to flux from neighboring sources than **psfFlux**. In addition to forced photometry PSF fluxes, a number of Boolean flags are also included in the **Forced Source catalog**.

The Forced Source catalog contains a total of 45 billion entries across the 790 million unique objects.

3.2.6. Forced Source at DIA Object

The Forced Source at DIA Object catalog is equivalent to the Forced Source catalog, but contains forced photometry measurements obtained at the positions of all the DIA objects in the DIA Object catalog.

The Forced Source at DIA Object catalog contains a total of 24 billion entries across the 230 million unique DIA objects.

3.2.7. DIA Object

The DIA Object catalog contains the astrophysical objects that DIA sources are associated with (i.e., the DIA objects). The DIA Object catalog contains only non-solar-system objects; solar system objects are, instead, recorded in the Solar System Object catalog (§3.2.9). When a DIA source is identified, the DIA Object and Solar System Object catalogs are searched for objects to associate it with. If no association is found, a new DIA object is created and the DIA source is associated to it. Along similar lines, an attempt has been made to associate DIA objects across multiple bands, meaning

²⁵ It is feasible for a marginal 5σ object in one counterfactual to be just below the 5σ detection threshold in another counterfactual.

the DIA Object catalog, like the Object catalog, associated DIA objects across multiple bands. By the nature of their detection, DIA objects are typically transient or variable. Therefore, the DIA Object catalog contains summary statistics of their fluxes, such as the mean and standard deviation over multiple epochs; users must refer to the Forced Source at DIA Object catalog (§3.2.6) or the DIA Source catalog for single-epoch flux measurements of DIA objects.

The DIA Object catalog contains data for 230 million DIA objects.

3.2.8. DIA Source

The DIA Source catalog contains data on all the sources detected at $> 5\sigma$ significance, including those associated with known solar system objects, in the difference images. Unlike sources detected in visit images, DIA sources detected in difference images have gone through an association step in which an attempt has been made to associate them into underlying objects called DIA objects. The DIA Source catalog consolidates all this information across multiple visits and bands. The detections reported in the DIA Source catalog have not undergone deblending.

The DIA Source catalog contains data for 1 billion DIA sources.

3.2.9. Solar System Object

The Solar System Object catalog and the MPC Orbit Database (MPCORB) carry information about solar system objects. The MPCORB table provides the MPC-computed orbital elements for all known asteroids, including those that Rubin discovered.

The Solar System Object catalog serves primarily to provide the mapping between the IAU designation of an object (listed in MPCORB), and the internal `ssObjectId` identifier, which is used as a key to find solar system object observations in the DIA Source and Solar System Source catalogs. The Solar System Object catalog contains data for 0.3 million solar system objects.

3.2.10. Solar System Source

The Solar System Source catalog contains data on all DIA sources that are either associated with previously known solar system objects, or have been confirmed as newly discovered solar system objects by confirmation of their orbital properties. As entries in the Solar System Source catalog stem from the DIA Source catalog, they have all been detected at $> 5\sigma$ significance in at least one band. The Solar System Source catalog contains data for 8.1 million solar system sources.

3.2.11. Calibration Reference Catalog

The Calibration Reference Catalog is used to perform astrometric and photometric calibration. It is a whole-sky catalog built specifically for LSST by combining data from multiple prior reference catalogs, since no single existing catalog had both the depth and coverage needed to calibrate LSST data. It contains only stellar sources. Full details on how the Calibration Reference Catalog was built are provided in (Ferguson et al. 2025).²⁶ We provide a brief summary here.

For the *grizy* bands, the input catalogs were (in order of decreasing priority): Dark Energy Survey (DES) Year 6 Calibration Stars (Rykoff et al. 2023); Gaia B or R Photometry (Gaia) (XP) synthetic magnitudes (Gaia Collaboration et al. 2023); the Panoramic Survey Telescope and Rapid Response System (Pan-STARRS) 3PI survey (Chambers et al. 2016); Data Release 2 of the SkyMapper survey (Onken et al. 2019); and Data Release 4 of the VLT Survey Telescope (VST) Asteroid Terrestrial-impact Last Alert System (ATLAS) survey (Shanks et al. 2015). For the *u* band, the input catalogs were (in order of decreasing priority): standard stars from Sloan Digital Sky Survey (SDSS) Data Release 16 (Ahumada et al. 2020); Gaia XP Synthetic Magnitudes (Gaia Collaboration et al. 2023); and synthetic magnitudes generated using Single Lens Reflex (SLR), which estimates the *u*-band flux from the *g*-band flux and *g* - *r* colors. These SLR estimates were used to boost the number of *u*-band reference sources, as otherwise the source density from the *u*-band input catalogs is too low to be useful for the LSST.

Throughout, the Calibration Reference Catalog uses the DES bandpasses for the *grizy* bands and the SDSS bandpass for the *u*-band; color transformations derived from high-quality sources were used to convert fluxes from the various input catalogs (some of which did not use the DES or SDSS bandpasses) to the respective bandpasses. All sources from the input catalogs are matched to Gaia Data Release 3 (DR3) sources for robust astrometric information, selecting only isolated sources (i.e., no neighbors within 1").

After collating the input catalogs and transforming the fluxes to the standard DES and SDSS bandpasses, the catalog was used to identify sources within a specific region of the sky. This process generated a set of standard columns containing positional and flux information, along with their associated uncertainties.

²⁶ In (Ferguson et al. 2025), the calibration reference catalog is referred to as “The Monster”.

3.3. Maps

3.3.1. Survey Property Maps

Survey Property maps summarize how properties such as observing conditions or exposure time vary across the observed sky. Each map provides the spatial distribution of a specific quantity at a defined sky position for each band by aggregating information from the images used to make the Deep Coadds. Maps are initially created per tract and then combined to produce a final consolidated map. At each sky location, represented by a spatial pixel in the [Hierarchical Equal-Area iso-Latitude Pixelisation \(HEALPix\)](#) (Górski et al. 2005) grid, values are derived using statistical operations, such as minimum, maximum, mean, weighted mean, or sum, depending on the property.

DP2 contains fourteen Survey Property maps. The available maps describe total exposure times, observation epochs (one each for the earliest, mean, and latest observation epoch), [Point Spread Function \(PSF\)](#) size and [shape](#) (one for each of the e^1 and e^2 shape parameters; see §??), PSF magnitude limits, sky background and noise levels, as well as astrometric shifts (one each for R.A. and decl.) and PSF distortions (one for each of the e^1 and e^2 shape parameters) due to wavelength-dependent atmospheric [Differential Chromatic Refraction \(DCR\)](#) effects.

They all use the dataset type format `deepCoadd_<PROPERTY>_consolidated_map_<STATISTIC>`. For example, `deepCoadd_exposure_time_consolidated_map_sum` provides a spatial map of the total exposure time accumulated per sky position in units of seconds. All maps are stored in [HealSparse](#)²⁷ format. Survey Property maps are only available via the Data Butler (§6.2.2).

3.3.2. HiPS Maps

[Hierarchical Progressive Survey \(IVOA standard\) \(HiPS\)](#) maps, offer an interactive way to explore seamless, multiband tiles of the sky regions covered by DP2, allowing for smooth panning and zooming. DP2 provides multiband [HiPS](#) images created by combining data from individual bands of Deep Coadd images, using an improved version (N. B. Lust et al. in preparation) of the algorithm presented in (Lupton et al. 2004). These images are false-color representations generated using various filter combinations for the red, green, and blue channels.

[HiPS](#) maps are only accessible through the [HiPS](#) viewer in the [RSP Portal \(§6.3\)](#) and cannot be accessed via the Data Butler (§6.2.2).

3.4. Ancillary Data Products

In addition to the science images and catalog data products described above, there are a number of ancillary data products included in DP2. While not science-level data products *per se*, their inclusion in DP2 will help to facilitate science.

3.4.1. Skymap

The DP2 skymap, `lsst_cells_v2`, is also included as a data product in DP2. The skymap data product can be used, for example, to identify the tract and patch that contains a given sky coordinate, or obtain the coordinates of the corners of a tract or patch.

A description of the `lsst_cells_v2`, in terms of the numbers and sizes of its tracts and patches, is given at the start of this section.

3.4.2. Standard Passbands

The corresponding transmission curves are provided as a DP2 data product. These datasets tabulate the full-system transmission of the six LSSTCam filters as a function of wavelength and were used as a reference for the LSSTCam DP2 photometry. The `standard_passband` dataset is keyed by band and is stored in Astropy table format.

3.4.3. Object Scarlet Models

The Object Scarlet Models dataset stores the scarlet models fit to each deblended detection in the Deep Coadd images. Each model captures an object's spatial profile (morphology) together with its per-band flux values, as assigned by the deblender when separating it from its blended neighbours. When combined with the corresponding Object catalog and Deep Coadd image, this dataset enables the reconstruction of the deblend model to see how a specific detection was deblended. There is one Object Scarlet Models dataset per patch that has coverage in DP2, and they are only available via the Data Butler.

3.4.4. Deep Coadd Input Summary

The Deep Coadd Input Summary table contains the visit and detector IDs of all the Visit Images that were selected and coadded to produce a given Deep Coadd image. Also included in the table are the relative weight of each Visit Image, inversely proportional to its variance, and the number of pixels it contributed to the Deep Coadd, since a given Visit Image may only partially cover a patch. There is one Deep Coadd Input

²⁷ A sparse [HEALPix](#) representation that efficiently encodes data values on the celestial sphere. <https://healsparse.readthedocs.io>

Summary per tract that has coverage in DP2, and they are only available via the Data Butler.

3.5. Metadata

3.5.1. Visit

The **Visit** catalog contains technical data for each visit collected at the time of observation, such as telescope pointing, camera rotation, airmass, exposure start and end time, and total exposure time. Some of the information contained within the **Visit** catalog is also contained in the Visit Detector catalog described below (e.g., exposure time), although the latter also includes information produced by the processing pipelines at the per-detector level, such as the PSF size and limiting magnitudes of a given Visit Image.

The Visit catalog contains 28698 entries.

3.5.2. Visit Detector

The Visit Detector catalog contains summary information for all Visit Images in DP2. In addition to technical information, such as the on-sky coordinates of the central pixel and measured pixel scale, the Visit Detector catalog contains a range of data quality measurements, such as whole-image summary statistics for the PSF size, zero-point and 5σ depth, sky background, sky noise, and the quality of astrometric solution. It provides an efficient method to access the properties of a Visit Image without needing to access the image data. When combined with the data contained in the Visit table described above, it provides a full picture of the telescope pointing and sky conditions at the time of observation.

The Visit Detector catalog contains 5.2 million entries.

3.5.3. Visit Summary

The Visit Summary table contains the same information as the Visit Detector catalog, but only for a single visit. There is one Visit Summary table per visit (meaning the Visit Detector table is the collation of all the Visit Summary tables), and they are only available via the Data Butler.

4. DATA RELEASE PROCESSING

5. PERFORMANCE CHARACTERIZATION AND KNOWN ISSUES

For information regarding the performance and known issues of the data contained in the Early DP2 release, we refer to the Data Quality, Image Artifacts, and Known Issues sections of the Release Documentation at <https://dp2.lsst.io>. A detailed analysis of the validation of DP2 will be provided at the time of its release (expected in the Oct–Dec 2026 range).

6. RUBIN SCIENCE PLATFORM

The **RSP** (Jurić et al. 2019) is a powerful, cloud-based environment for scientific research and analysis of petascale astronomical survey data and serves as the primary interface for scientists to access, visualize, and conduct next-to-the-data analysis of Rubin and **LSST** data. Built around a “bring the compute to the data” principle, it eliminates the need for users to download massive datasets, instead providing Rubin data rights holders with secure, single-sign-on access to data and services via three distinct user-facing Aspects: a *Portal*, which facilitate interactive exploration of the data; a JupyterLab-based *Notebook* environment for data analysis using Python; and an extensive set of *APIs* that enable programmatic access to both data and services.

The three Aspects are designed to be fully integrated, enabling seamless workflows across the **RSP**. The **RSP** is supported by a number of backend services, including databases, files, and batch computing. Support for collaborative work through shared workspaces is also included in the **RSP**. The data products described in §3 are accessible via all three Aspects, and the system facilitates operations such as starting a query in one Aspect and retrieving its results in another. Figure 1 shows the **RSP** landing page in the Google cloud.




Figure 1. The **RSP** landing page at <https://data.lsst.cloud/> showing the three user-facing Aspects as well as links to documentation and support information.

A preview of the **RSP** was first launched on Google Cloud in 2022, operating under a shared-risk model to support DP0 (O’Mullane et al. 2024b). DP1 built on that early preview to bring major upgrades including the hybrid on-premises **USDF**/cloud deployment, a broader set of **International Virtual Observatory Alliance (IVOA)** API services (SIAv2, SODA cutouts, HiPS, WebDAV, alongside the existing **Table Access**

Protocol (IVOA standard) (TAP)), and a fully integrated, read-only Data Butler (S3 store + AlloyDB registry) across all three RSP Aspects.

The size (3.5 TB) and complexity of DP1 did not, on their own, necessitate the full functionality of the RSP for analysis; a traditional “download the data” model would still have sufficed. However, it was important to provide users the same experience and interfaces they will rely on for later, larger releases, and to get them started in this new paradigm on a small, manageable dataset. DP2, by contrast, marks a new regime, with its size of 59.15 TB outside the range of the traditional download model. The experience users gained working with DP1 through the RSP interfaces can now be applied directly to DP2. The release of DP2 builds on this evolving platform to bring additional upgrades and new functionality, including expanded user query-history and query-status capabilities, context-aware in-application documentation (e.g., contextual help linked directly from Portal table columns), and support for the ADQL spatial INTERSECTS() operation.

The RSP remains under active development, with incremental improvements being rolled out as they mature. During the Rubin Early Science Phase, the RSP will continue to operate under a shared-risk model. This section outlines the RSP functionality available at the time of the DP2 release and provides an overview of planned future capabilities.

6.1. Rubin US Data Access Center

The Rubin US DAC utilizes a novel hybrid architecture that combines on-premises infrastructure at the USDF at SLAC with flexible, scalable resources in the Google cloud (O’Mullane et al. 2024a). User-facing services are deployed in the cloud to support dynamic scaling in response to user demand and to simplify the provisioning and management of large numbers of science user accounts, whereas the majority of the static data products described in §3 are stored on premises at the USDF, benefiting from cost-effective mass storage and close integration with Rubin’s data processing infrastructure, also located at the USDF. For all imaging data except coadds, the Data Butler (§6.2.2) provides the interface between cloud-based users and the on-premises data; for catalog data, a cloud-based TAP client (§6.2.1) submits queries to the on-premises Qserv database cluster (§6.5) and retrieves the results. Deep Coadd images are the exception to this hybrid model: because there is one coadd per sky region per band, a relatively small and size-bounded dataset, they are hosted directly in Google Cloud to benefit from cloud-native access.

DP2 is being delivered in two phases because full end-to-end integration of the hybrid RSP for images was not yet complete by the time data processing (§??) finished; the hybrid model supporting Qserv (§??) catalog access over TAP, by contrast, is already in production for EDP2. Rubin therefore adopted a phased release, given that neither coadds nor catalogs depend on the image-side hybrid RSP integration readiness. Early DP2 (EDP2), comprising the deep coadd images and all catalogs, was released on 2026-07-27, while the full DP2 release, adding the visit-level data products once hybrid image integration is demonstrated, is expected in the Oct–Dec 2026 range.

6.2. Application Programming Interface Aspect

The Application Programming Interface (API) Aspect provides a comprehensive set of user-facing interfaces for programmatic access to the DP1 data products, through both IVOA-compliant services and the Rubin Data Butler. IVOA services enable standard queries and integration with existing tools, while the Butler facilitates advanced data processing with the LSST Science Pipelines.

Since the DP1 release, additional services have been added as the RSP moves toward full functionality, but some planned services remain unavailable, and certain data products are only accessible via the Butler. This section provides an overview of the available IVOA services and Butler access.

6.2.1. IVOA Services

Rubin has adopted a Virtual Observatory (VO)-first design philosophy, prioritizing compliance with IVOA standard interfaces to foster interoperability, standardization, and collaboration. In cases where standardized protocols have yet to be established, additional services have been introduced to complement these efforts. This approach ensures that the RSP can be seamlessly integrated with community-standard tools such as Tool for Operations on Catalogues And Tables (TOPCAT) (Taylor 2011) and Aladin (Bonnarel et al. 2000; Boch & Fernique 2014; Baumann et al. 2022), as well as libraries such as PyVO (Graham et al. 2014).

The user-facing APIs are also used internally within the RSP, creating a unified design that ensures consistent and reproducible workflows across all three Aspects. This reduces code duplication, simplifies maintenance, and ensures all users, both internal and external, access data in the same way. For example, an Astronomical Data Query Language (IVOA standard) (ADQL) query on the Object catalog via TAP yields identical results whether run from the Portal, Notebook, or an external client.

Since the [DP1](#) release, the RSP has continued to progress towards a full-featured platform, with a number of new services and capabilities added in support of [DP2](#). The following [IVOA](#) services are available at the time of the [DP2](#) release.

- **TAP service:** A TAP service ([Dowler et al. 2019](#)) enables queries of catalog data via the [IVOA](#)-standard [ADQL](#), a dialect of SQL92 with spherical geometry extensions. The main TAP service for [DP2](#) runs on the Rubin-developed [Qserv](#) database (§ 6.5), which hosts the core science tables described in §??, as well as the [Visit](#) database. It also provides image metadata in the [IVOA ObsCore](#) format via the standard `ivoa.ObsCore` table, making it an “ObsTAP” service ([Louys et al. 2017](#)). The TAP service is based on the [Canadian Astronomy Data Centre \(CADC\)](#)’s open-source Java TAP implementation²⁸, modified for the exact query language accepted by [Qserv](#). Since [DP1](#), [ADQL](#) support has been expanded to include the spatial `INTERSECTS()` operation, and temporary-table uploads now provide multi-object query capabilities across both [Qserv](#)- and [Postgres](#)-backed catalogs. It currently supports a large subset of [ADQL](#), with limitations documented in the Data Release materials (see §??) and exposed via the TAP capabilities endpoint where possible.

The TAP service provides metadata annotations consistent with the standard, including table and column descriptions, indications of foreign-key relationships between tables, and column metadata such as units and [IVOA](#) unified content descriptors (UCDs).

- **Image access services:** Rubin image access services are compliant with [IVOA](#) Simple Image Access Protocol, version 2 (SIAv2; [Jenness et al. 2024](#); [Dowler et al. 2015](#)) for discovering and accessing astronomical images based on [metadata](#). SIAv2 is a [REpresentational State Transfer \(REST\)](#)-based protocol designed for the discovery and retrieval of image data. It allows, for instance, querying all images in a given band over a defined sky region and time period.

Users identify an image or observation of interest and query the service. The result set includes [metadata](#) about the image, such as the sky position, time, or band, and a data-access URL, which

includes an [IVOA](#) Identifier uniquely identifying the dataset ([Jenness & Dubois-Felsmann 2026](#)), allowing the dataset to be retrieved or a cutout requested via [Server-side Operations for Data Access \(IVOA standard\) \(SODA\)](#) ([SODA](#); [Bonnarel et al. 2017](#)).

- **Image Cutout Service:** The Rubin Cutout Service ([Allbery 2023, 2024](#)) is based on the [IVOA SODA](#) standard ([Bonnarel et al. 2017](#)). Users submit requests specifying sky coordinates and the cutout size as the radius from the coordinates, and the service performs the operation on the full image and returns a result set. As with [DP1](#), the Cutout Service for [DP2](#) remains a single-cutout service, where N cutout requests require N independent synchronous calls; a bulk Cutout Service, ([Jenness et al. 2026](#)) is now planned for [Data Release 1 \(DR1\)](#).
- **HiPS Data Service:** An authenticated [HiPS](#) ([Fernique et al. 2017](#)) Data Service for seamless pan-and-zoom access to large-scale coadds. It supports fast interactive progressive image exploration at a range of resolutions.
- **Web Distributed Authoring and Versioning (WebDav):** A [WebDav](#) service is provided to enable users to remotely manage, edit, and organize files and directories on the [RSP](#) as if they were local files on their own computer. This is especially useful for local development.

Beyond these [IVOA](#) services, [DP2](#) introduces expanded user query-history and query-status capabilities, initial context-aware, in-application documentation (e.g., contextual help linked directly from Portal table columns).

Looking ahead to [DR1](#), planned RSP functionality includes bulk cutout services (both for individual targets across a range of epochs and for lists of multiple targets), a [PSF](#) retrieval service, an on-demand forced-photometry service, data product recreation services, parallel and batch computing, support for persistent user-created databases and their federation with project-provided databases, support for collaborative work, and full LSSTCam focal-plane visualizations. Access to GPUs and the ability to bring additional user-funded compute resources to the RSP are under consideration for the post-[DR1](#) period.

6.2.2. Data Butler

The Rubin Data Butler ([Jenness et al. 2022](#); [Lust et al. 2023](#)), is a high-level interface designed to facili-

²⁸ <https://github.com/opencadc/tap>

tate seamless access to data for both users and software systems. This includes managing storage formats, physical locations, data staging, and database mappings. A [Butler](#) repository contains two components.

- The *Data Store*. A physical storage system for datasets, e.g., a [Portable Operating System Interface \(POSIX\)](#) file system or [\(Amazon\) Simple Storage Service \(S3\)](#) object store.
- The *Registry*. A [Structured Query Language \(SQL\)](#)-compatible database that stores metadata about the datasets in the Data Store.

In the context of the [Butler](#), a dataset refers to a unique data product, such as an image, catalog or map, generated by the observatory or processing pipelines. Datasets belong to one of the various types of data products, described in §3. The [Butler](#) ensures that each dataset is uniquely identifiable by a combination of three pieces of information: a data coordinate, a dataset type, and a run collection.

For a DP2 deep coadd on a [patch](#) of sky, the dimensions would be the [tract](#), [patch](#) and band, specified as a `dataId`, e.g., `dataId={'tract': 11399, 'patch': 16, 'band': 'g'}` and is associated with the `deep_coadd` `DatasetType`. In contrast, the `visit_table` and `visit_detector_table` catalogs described in §3 have only the instrument dimension, since each is a single table covering all visits (or all visit-detector combinations) for a given instrument; retrieving either therefore requires only `dataId={'instrument': 'LSSTCam'}`. The `visit_summary` catalog, by contrast, is keyed on band, instrument, `day_obs`, `physical_filter`, and visit, e.g., `dataId={'instrument': 'LSSTCam', 'visit': 2025071100074, 'band': 'g', 'physical_filter': 'g_6', 'day_obs': 20250711}`, and contains one row per detector for that visit.

As for DP1, the [EDP2](#) [Butler](#) repository containing the catalogs and coadds described in §3 is hosted in the Google Cloud using an [S3](#)-compatible store for datasets and AlloyDB, a PostgreSQL-compatible database for the Registry.

6.2.3. Remote Programmatic Access

The Rubin [RSP API](#) can be accessed remotely by data rights holders by creating a user security token. This token can then be used as a bearer token for [API](#) calls to the [RSP](#) services. This capability is especially useful for remote data analysis using tools such as [TOPCAT](#), and for enabling third-party systems, such as community Alert brokers, to access Rubin data. Additionally, it supports remote development, allowing for more flexible workflows and integration with external systems.

6.3. Portal Aspect

The Portal Aspect provides an interactive web-based environment for exploratory data discovery, filtering, querying, and visualization of both image and catalog data, without requiring programming expertise. It enables users to access and analyze large datasets via tools for catalog queries, image browsing, time-series inspection, and crossmatching.

The Portal is built on [Firefly](#) ([Wu et al. 2019](#)), a web application framework developed by the Infrared Processing and Analysis Center. [Firefly](#) provides interactive capabilities such as customizable table views, image overlays, multipanel visualizations, and synchronized displays linking catalog and image data.

Designed to support both exploratory data access and detailed scientific investigation, the Portal delivers an intuitive user experience, allowing users to visually analyze data while retaining access to underlying metadata and query controls.

6.4. Notebook Aspect

The Notebook Aspect provides an interactive, web-based environment built on Jupyter Notebooks, enabling users to write and execute Python code directly on Rubin and [LSST](#) data without downloading it locally. It offers programmatic access to Rubin and [LSST](#) data products, allowing users to query and retrieve datasets, manipulate and display images, compute derived properties, plot results, and reprocess data using the [LSST Science Pipelines](#) (§??). The environment comes preinstalled with the Pipelines and a broad set of widely used astronomical [software](#) tools, supporting immediate and flexible data analysis.

6.5. Databases

The user-facing Aspects of the [RSP](#) are supported by several backend databases that store catalog data products, image metadata, and other derived datasets. The [schema](#) for DP2 and other Rubin databases are available online at <https://sdm-schemas.lsst.io/dp2.html>.

6.5.1. Qserv

The final 10 yr [LSST](#) catalog is expected to reach 15 PB and contain measurements for billions of stars and galaxies across trillions of detections. To support efficient storage, querying, and analysis of this dataset, Rubin Observatory developed Qserv ([Wang et al. 2011](#); [Mueller et al. 2025](#)), a scalable, parallel, distributed SQL database system. [Qserv](#) partitions data over approximately equal-area regions of the celestial sphere, replicates data to ensure resilience and high availability, and uses shared scanning to reduce overall I/O load. It

also supports a package of scientific user-defined functions (SciSQL²⁹) simplifying complex queries involving spherical geometry, statistics, and photometry. Qserv is built on robust production-quality components, including MariaDB³⁰ and XRootD³¹. Qserv runs at the USDF and user access to catalog data is via the TAP service (§6.2.1). This enables catalog-based analysis through both the RSP Portal and Notebook Aspects.

While DP1 was small enough that Qserv’s full capabilities were not strictly required, DP2, at 59.15 TB, is substantially larger and begins to exercise Qserv at a scale meaningfully closer to its intended operating regime, though still well short of the petabyte-scale catalogs expected for the 10 yr LSST survey.

Qserv is open-source and available on GitHub at <https://github.com/lsst/qserv>.

6.5.2. OpSim

DP2 provides an opsim-format SQLite database of visit-level observational metadata for the visits included in DP2, downloadable from https://s3df.slac.stanford.edu/data/rubin/sim-data/prelsst/dp2_visits.db, with additional documentation available at https://survey-strategy.lsst.io/progress/prelsst/prelsst_dp2.html#visit-database. Like the Butler `visit_table`, this database contains visit-level entries, but it has no equivalent of the `visit_detector_table` described in §?? and so provides no per-detector information. Its quantities (e.g., PSF FWHM, 5σ point-source depth) are drawn from the real-time Rapid Analysis QuickLook processing, () performed on the summit rather than from Data Release Processing (DRP) processing. Values may therefore differ somewhat from the corresponding DRP-derived quantities in the Butler `visit_table`. The database also includes some quantities not otherwise available from DRP outputs, such as AOS Zernike coefficients, as well as predicted quantities useful for comparison against simulated survey performance.

7. SUPPORT FOR COMMUNITY SCIENCE

Rubin Observatory’s model to support its large, diverse science community to access and analyze DP2 focuses on providing documentation, tutorials, and live virtual learning opportunities alongside an open platform for issue reporting and resolution that enables crowd-sourced solutions. These support resources are

all available via the main hub of Rubin’s “For Scientists” website³².

The data release documentation for DP2³³ provides an overview of the LSSTCam observations, data quality summary plots, and descriptions of the contents and access methods for all catalog and image datasets. All tutorials (NSF-DOE Vera C. Rubin Observatory 2021) that demonstrate DP2 access via the Rubin Science Platform are delivered in the data release documentation, and executable versions of the Jupyter notebook tutorials are also available from a drop-down menu within the Notebook Aspect.

Early DP2 was released on the first day of the annual Rubin Community Workshop, July 27 2026. A dedicated “Get started with DP2” session was added to the Wednesday agenda. Four DP2-onboarding virtual sessions were run in the first week of August, a mix of presentations, tutorials, and time for issue resolution.

The Rubin Community Forum³⁴ is the open platform for user issue reporting and resolution. Rubin staff monitor all new topics posted in the Support category and ensure that all requests for help get a response and a solution. The open nature of the platform also enables peer-to-peer support that avoids knowledge bottlenecks and hastens the time to resolution.

8. SUMMARY AND FUTURE RELEASES

DP2 provides the first release of science-grade data products from the LSSTCam focal plane. As for DP1, DP2 is part of Rubin’s Early Science Program (Guy et al. 2026) and carries a 2 yr proprietary period during which access is restricted to LSST data rights holders via the cloud-based RSP.

DP2 is delivered in two phases: Early DP2 (EDP2), which provides deep coadded images, all science catalogs, and ancillary data products including throughput curves, survey property maps, and calibration products; and DP2, planned for late 2026, which adds single-epoch visit-level and difference-image data products to complete the release. The coadded images and catalogs delivered in EDP2 will not be reprocessed for DP2; the data are unchanged, and DP2 supersedes and fully encompasses EDP2 by adding the remaining products to the same release. Analyses begun on EDP2 products require no modification.

DP2 represents a substantial advance relative to DP1: the full LSSTCam focal plane replaces the single-raft LSSTComCam, all six *ugrizy* filters are available in

²⁹ <https://smonkewitz.github.io/scisql/>

³⁰ <https://www.mariadb.org/>

³¹ <https://xrootd.org/>

³² <https://rubinobservatory.org/for-scientists>

³³ <https://dp2.lsst.io>

³⁴ <https://community.lsst.org/>

place of three, and the sky area is larger by approximately two orders of magnitude. Nevertheless, DP2 reflects the state of an observatory still in the process of optimization: delivered image quality had not yet consistently reached design specification, template coverage is incomplete across the full sky, and several pipeline components remain under active development. Users are encouraged to consult the release documentation at <https://dp2.lsst.io> and the known issues register maintained there before beginning scientific analyses.

The body of science-grade LSSTCam data acquired between the end of the DP2 observing window on 2026-01-06 and the formal start of the LSST on 2026-06-29 will be processed and made available to LSST data rights holders through the same RSP services. The format and timing of this release are still under discussion; planning will be finalised once DP2 has been delivered.

The first LSST Data Release, DR1, will be derived from a reprocessing of approximately the first year of LSST survey data and is currently expected in mid-2028; its precise definition remains under consideration. DR1 will mark a qualitative advance beyond DP2: observations collected under the community-developed baseline survey strategy (Rubin’s Survey Cadence Optimization Committee et al. 2022, 2023, 2025; The Rubin Observatory Survey Cadence Optimization Committee 2026) will yield substantially broader and more uniform sky coverage across the full Wide-Fast-Deep (WFD) footprint of 18000 deg², a balanced six-band filter distribution, and improved delivered image quality as AOS and dome thermal management continue to mature (NSF-DOE Vera C. Rubin Observatory 2026). Photometric and astrometric calibration will benefit from a full year of all-sky coverage, difference-imaging templates will span a far greater fraction of the footprint and processing will draw on a more mature version of the LSST Science Pipelines refined through DP2 experience and community engagement. Subsequent LSST Data Releases are planned at an approximately yearly cadence,

each a complete reprocessing of all survey data accumulated to that point, with progressively deeper coadds and improved calibration. Plans will be updated through revisions to Guy et al. (2026) and on the Rubin community forum at <https://community.lsst.org>.

ACKNOWLEDGMENTS

. This material is based upon work supported in part by the National Science Foundation through Cooperative Agreements AST-1258333 and AST-2241526 and Cooperative Support Agreements AST-1202910 and AST-2211468 managed by the Association of Universities for Research in Astronomy (AURA), and the Department of Energy under Contract No. DE-AC02-76SF00515 with the SLAC National Accelerator Laboratory managed by Stanford University. Additional Rubin Observatory funding comes from private donations, grants to universities, and in-kind support from LSST-DA Institutional Members.

This work has been supported by the French National Institute of Nuclear and Particle Physics (IN2P3) through dedicated funding provided by the National Center for Scientific Research (CNRS).

This work has been supported by STFC funding for UK participation in LSST, through grant ST/Y00292X/1.

Facilities: Rubin:Simonyi (LSSTComCam), Rubin:USDAC

Software: Rubin Data Butler (Jenness et al. 2022), LSST Science Pipelines (Rubin Observatory Science Pipelines Developers 2025), LSST Feature Based Scheduler v3.0 (Jones 2026) Astropy (Astropy Collaboration et al. 2013, 2018, 2022) PIFF (Jarvis et al. 2021), GBDES (Bernstein 2022), Qserv (Wang et al. 2011; Mueller et al. 2025), Slurm, HTCondor, CVMFS, FTS3, ESNet

APPENDIX

REFERENCES

- Ahumada, R., Allende Prieto, C., Almeida, A., et al. 2020, ApJS, 249, 3, doi: [10.3847/1538-4365/ab929e](https://doi.org/10.3847/1538-4365/ab929e)
- Allbery, R. 2023, IVOA SODA implementation experience, SQuaRE Technical Note SQR-063, NSF-DOE Vera C. Rubin Observatory. <https://sqr-063.lsst.io/>
- . 2024, Draft IVOA SODA web service specification, SQuaRE Technical Note SQR-093, NSF-DOE Vera C. Rubin Observatory. <https://sqr-093.lsst.io/>
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, A&A, 558, A33, doi: [10.1051/0004-6361/201322068](https://doi.org/10.1051/0004-6361/201322068)

- Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, 156, 123, doi: [10.3847/1538-3881/aabc4f](https://doi.org/10.3847/1538-3881/aabc4f)
- Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, *ApJ*, 935, 167, doi: [10.3847/1538-4357/ac7c74](https://doi.org/10.3847/1538-4357/ac7c74)
- Baumann, M., Boch, T., Pineau, F.-X., et al. 2022, in *Astronomical Society of the Pacific Conference Series*, Vol. 532, *Astronomical Data Analysis Software and Systems XXX*, ed. J. E. Ruiz, F. Pierfederici, & P. Teuben, 7
- Bernstein, G., Gruendl, R., & Kent, S. 2016, Limiting Magnitude, t_{eff} , and Image Quality in DES Year 1, Tech. rep., Fermi National Accelerator Laboratory (FNAL), Batavia, IL (United States), doi: [10.2172/1250877](https://doi.org/10.2172/1250877)
- Bernstein, G. M. 2022, gbdes: DECam instrumental signature fitting and processing programs, *Astrophysics Source Code Library*, record ascl:2210.011. <http://ascl.net/2210.011>
- Bianco, F. B., Ivezić, Ž., Jones, R. L., et al. 2022, *ApJS*, 258, 1, doi: [10.3847/1538-4365/ac3e72](https://doi.org/10.3847/1538-4365/ac3e72)
- Blum, R., & the Rubin Operations Team. 2020, Vera C. Rubin Observatory Data Policy, Data Management Operations Controlled Document RDO-013, NSF-DOE Vera C. Rubin Observatory. <https://ls.st/RDO-013>
- Boch, T., & Fernique, P. 2014, in *Astronomical Society of the Pacific Conference Series*, Vol. 485, *Astronomical Data Analysis Software and Systems XXIII*, ed. N. Manset & P. Forshay, 277
- Bonnarel, F., Dowler, P., Demleitner, M., Tody, D., & Dempsey, J. 2017, IVOA Server-side Operations for Data Access Version 1.0, IVOA Recommendation 17 May 2017, doi: [10.5479/ADS/bib/2017ivoa.spec.0517B](https://doi.org/10.5479/ADS/bib/2017ivoa.spec.0517B)
- Bonnarel, F., Fernique, P., Bienaymé, O., et al. 2000, *A&AS*, 143, 33, doi: [10.1051/aas:2000331](https://doi.org/10.1051/aas:2000331)
- Bosch, J. F. 2026, Challenges serializing Rubin Observatory data products, Data Management Technical Note DMTN-339, NSF-DOE Vera C. Rubin Observatory. <https://dmtn-339.lsst.io/>
- Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2016, arXiv e-prints, arXiv:1612.05560, doi: [10.48550/arXiv.1612.05560](https://doi.org/10.48550/arXiv.1612.05560)
- Dowler, P., Bonnarel, F., & Tody, D. 2015, IVOA Simple Image Access Version 2.0, IVOA Recommendation 23 December 2015, doi: [10.5479/ADS/bib/2015ivoa.spec.1223D](https://doi.org/10.5479/ADS/bib/2015ivoa.spec.1223D)
- Dowler, P., Rixon, G., Tody, D., & Demleitner, M. 2019, Table Access Protocol Version 1.1, IVOA Recommendation 27 September 2019, doi: [10.5479/ADS/bib/2019ivoa.spec.0927D](https://doi.org/10.5479/ADS/bib/2019ivoa.spec.0927D)
- Ferguson, P. S., Rykoff, E. S., Carlin, J. L., Saunders, C., & Parejko, J. K. 2025, The Monster: A reference catalog with synthetic ugrizy-band fluxes for the Vera C. Rubin observatory, Data Management Technical Note DMTN-277, NSF-DOE Vera C. Rubin Observatory, doi: [10.71929/rubin/2583688](https://doi.org/10.71929/rubin/2583688)
- Fernique, P., Allen, M., Boch, T., et al. 2017, HiPS - Hierarchical Progressive Survey Version 1.0, IVOA Recommendation 19 May 2017, doi: [10.5479/ADS/bib/2017ivoa.spec.0519F](https://doi.org/10.5479/ADS/bib/2017ivoa.spec.0519F)
- Gaia Collaboration, Montegriffo, P., Bellazzini, M., et al. 2023, *A&A*, 674, A33, doi: [10.1051/0004-6361/202243709](https://doi.org/10.1051/0004-6361/202243709)
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, 622, 759, doi: [10.1086/427976](https://doi.org/10.1086/427976)
- Graham, M., Plante, R., Tody, D., & Fitzpatrick, M. 2014, PyVO: Python access to the Virtual Observatory, *Astrophysics Source Code Library*, record ascl:1402.004
- Guy, L. P., AlSayyad, Y., Bechtol, K., et al. 2026, Rubin Observatory Plans for an Early Science Program, Technical Note RTN-011, NSF-DOE Vera C. Rubin Observatory, doi: [10.71929/rubin/2584021](https://doi.org/10.71929/rubin/2584021)
- Howard, J., Reil, K., Claver, C., et al. 2018, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 10700, *Ground-based and Airborne Telescopes VII*, ed. H. K. Marshall & J. Spyromilio, 107003D, doi: [10.1117/12.2312684](https://doi.org/10.1117/12.2312684)
- Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, 873, 111, doi: [10.3847/1538-4357/ab042c](https://doi.org/10.3847/1538-4357/ab042c)
- Jarvis, M., Bernstein, G. M., Amon, A., et al. 2021, *MNRAS*, 501, 1282, doi: [10.1093/mnras/staa3679](https://doi.org/10.1093/mnras/staa3679)
- Jenness, T., & Dubois-Felsmann, G. P. 2026, IVOA Identifier Usage at the Rubin Observatory, Data Management Technical Note DMTN-302, NSF-DOE Vera C. Rubin Observatory, doi: [10.71929/rubin/2583848](https://doi.org/10.71929/rubin/2583848)
- Jenness, T., Dubois-Felsmann, G. P., & Bosch, J. F. 2026, Bulk Cutout Service Implementation Options, Data Management Technical Note DMTN-326, NSF-DOE Vera C. Rubin Observatory. <https://dmtn-326.lsst.io/>
- Jenness, T., Voutsinas, S., Dubois-Felsmann, G. P., & Salnikov, A. 2024, arXiv e-prints, arXiv:2501.00544, doi: [10.48550/arXiv.2501.00544](https://doi.org/10.48550/arXiv.2501.00544)
- Jenness, T., Bosch, J. F., Salnikov, A., et al. 2022, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 12189, *Software and Cyberinfrastructure for Astronomy VII*, 1218911, doi: [10.1117/12.2629569](https://doi.org/10.1117/12.2629569)
- Jones, R. L. 2026, Scheduling the LSST: Commissioning the Rubin Feature Based Scheduler, Simulations Team Technical Note SMTN-020, NSF-DOE Vera C. Rubin Observatory. <https://smtn-020.lsst.io/>

- Jurić, M., Ciardi, D., Dubois-Felsmann, G., & Guy, L. 2019, LSST Science Platform Vision Document, Systems Engineering Controlled Document LSE-319, NSF-DOE Vera C. Rubin Observatory, doi: [10.71929/rubin/2587242](https://doi.org/10.71929/rubin/2587242)
- Louys, M., Tody, D., Dowler, P., et al. 2017, Observation Data Model Core Components, its Implementation in the Table Access Protocol Version 1.1, IVOA Recommendation 09 May 2017, doi: [10.5479/ADS/bib/2017ivoa.spec.0509L](https://doi.org/10.5479/ADS/bib/2017ivoa.spec.0509L)
- LSST Dark Energy Science Collaboration (LSST DESC), Abolfathi, B., Alonso, D., et al. 2021, ApJS, 253, 31, doi: [10.3847/1538-4365/abd62c](https://doi.org/10.3847/1538-4365/abd62c)
- Lupton, R., Blanton, M. R., Fekete, G., et al. 2004, PASP, 116, 133, doi: [10.1086/382245](https://doi.org/10.1086/382245)
- Lust, N. B., Jenness, T., Bosch, J. F., et al. 2023, arXiv e-prints, arXiv:2303.03313, doi: [10.48550/arXiv.2303.03313](https://doi.org/10.48550/arXiv.2303.03313)
- Mueller, F., Gaponenko, I., Gates, J., et al. 2025, in Astronomical Society of the Pacific Conference Series, Vol. 538, Astronomical Data Analysis Software and Systems XXXII, ed. S. Gaudet, D. Bohlender, S. Gwyn, A. Hincks, & P. Teuben, 114, doi: [10.26624/XCPI7375](https://doi.org/10.26624/XCPI7375)
- NSF-DOE Vera C. Rubin Observatory. 2021, Rubin Observatory LSST Tutorials [Computer Software], NSF-DOE Vera C. Rubin Observatory, doi: [10.11578/rubin/dc.20250909.20](https://doi.org/10.11578/rubin/dc.20250909.20)
- . 2025, Legacy Survey of Space and Time Data Preview 1 [Data set], NSF-DOE Vera C. Rubin Observatory, doi: [10.71929/RUBIN/2570308](https://doi.org/10.71929/RUBIN/2570308)
- . 2026, Image Quality Improvement Plan, Summit Operations Technical Note SOTN-004, NSF-DOE Vera C. Rubin Observatory. <https://sotn-004.lsst.io/>
- Oke, J. B., & Gunn, J. E. 1983, ApJ, 266, 713, doi: [10.1086/160817](https://doi.org/10.1086/160817)
- O'Mullane, W., AlSayyad, Y., Chiang, J., et al. 2024a, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 13101, Software and Cyberinfrastructure for Astronomy VIII, ed. J. Ibsen & G. Chiozzi, 131012B, doi: [10.1117/12.3018005](https://doi.org/10.1117/12.3018005)
- O'Mullane, W., Economou, F., Huang, F., et al. 2024b, in Astronomical Society of the Pacific Conference Series, Vol. 535, Astronomical Data Analysis Software and Systems XXXI, ed. B. V. Hugo, R. Van Rooyen, & O. M. Smirnov, 227, doi: [10.48550/arXiv.2111.15030](https://doi.org/10.48550/arXiv.2111.15030)
- Onken, C. A., Wolf, C., Bessell, M. S., et al. 2019, PASA, 36, e033, doi: [10.1017/pasa.2019.27](https://doi.org/10.1017/pasa.2019.27)
- Petrosian, V. 1976, ApJL, 210, L53, doi: [10.1086/18230110.1086/182253](https://doi.org/10.1086/18230110.1086/182253)
- Rubin, V. C., & Ford, Jr., W. K. 1970, ApJ, 159, 379, doi: [10.1086/150317](https://doi.org/10.1086/150317)
- Rubin, V. C., Ford, Jr., W. K., & Thonnard, N. 1980, ApJ, 238, 471, doi: [10.1086/158003](https://doi.org/10.1086/158003)
- Rubin Observatory Science Pipelines Developers. 2025, The LSST Science Pipelines Software: Optical Survey Pipeline Reduction and Analysis Environment, Project Science Technical Note PSTN-019, NSF-DOE Vera C. Rubin Observatory, doi: [10.71929/rubin/2570545](https://doi.org/10.71929/rubin/2570545)
- Rubin's Survey Cadence Optimization Committee, Bauer, F. E., Brough, S., et al. 2022, Survey Cadence Optimization Committee's Phase 1 Recommendation, Project Science Technical Note PSTN-053, NSF-DOE Vera C. Rubin Observatory, doi: [10.71929/rubin/2584276](https://doi.org/10.71929/rubin/2584276)
- Rubin's Survey Cadence Optimization Committee, Bauer, F. E., Bianco, F. B., et al. 2023, Survey Cadence Optimization Committee's Phase 2 Recommendations, Project Science Technical Note PSTN-055, NSF-DOE Vera C. Rubin Observatory, doi: [10.71929/rubin/2585249](https://doi.org/10.71929/rubin/2585249)
- Rubin's Survey Cadence Optimization Committee, Bianco, F. B., Jones, R. L., et al. 2025, Survey Cadence Optimization Committee's Phase 3 Recommendations, Project Science Technical Note PSTN-056, NSF-DOE Vera C. Rubin Observatory, doi: [10.71929/rubin/2585402](https://doi.org/10.71929/rubin/2585402)
- Rykoff, E. S., Tucker, D. L., Burke, D. L., et al. 2023, arXiv e-prints, arXiv:2305.01695, doi: [10.48550/arXiv.2305.01695](https://doi.org/10.48550/arXiv.2305.01695)
- Shanks, T., Metcalfe, N., Chehade, B., et al. 2015, MNRAS, 451, 4238, doi: [10.1093/mnras/stv1130](https://doi.org/10.1093/mnras/stv1130)
- SLAC National Accelerator Laboratory, & NSF-DOE Vera C. Rubin Observatory. 2024, LSST Commissioning Camera, SLAC National Accelerator Laboratory (SLAC), Menlo Park, CA (United States), doi: [10.71929/RUBIN/2561361](https://doi.org/10.71929/RUBIN/2561361)
- Stalder, B., Reil, K., Claver, C., et al. 2020, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 11447, Ground-based and Airborne Instrumentation for Astronomy VIII, ed. C. J. Evans, J. J. Bryant, & K. Motohara, 114470L, doi: [10.1117/12.2561132](https://doi.org/10.1117/12.2561132)
- Stalder, B., Reil, K., Aguilar, C., et al. 2022, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 12184, Ground-based and Airborne Instrumentation for Astronomy IX, ed. C. J. Evans, J. J. Bryant, & K. Motohara, 121840J, doi: [10.1117/12.2630184](https://doi.org/10.1117/12.2630184)
- Stalder, B., Munoz, F., Aguilar, C., et al. 2024, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 13094, Ground-based and Airborne Telescopes X, ed. H. K. Marshall, J. Spyromilio, & T. Usuda, 1309409, doi: [10.1117/12.3019266](https://doi.org/10.1117/12.3019266)

- Taylor, M. 2011, TOPCAT: Tool for Operations on Catalogues And Tables, Astrophysics Source Code Library, record ascl:1101.010
- The Rubin Observatory Survey Cadence Optimization Committee. 2026, Survey Cadence Optimization Committee's Survey Start Recommendations, Project Science Technical Note PSTN-057, NSF-DOE Vera C. Rubin Observatory. <https://pstn-057.lsst.io/>
- Wang, D. L., Monkewitz, S. M., Lim, K.-T., & Becla, J. 2011, in State of the Practice Reports, SC '11 (New York, NY, USA: ACM), 12:1–12:11, doi: [10.1145/2063348.2063364](https://doi.org/10.1145/2063348.2063364)
- Wu, X., Roby, W., Goldian, T., et al. 2019, in Astronomical Society of the Pacific Conference Series, Vol. 521, Astronomical Data Analysis Software and Systems XXVI, ed. M. Molinaro, K. Shortridge, & F. Pasian, 32

Glossary

ADQL: Astronomical Data Query Language (IVOA standard).

AOS: Active Optics System.

API: Application Programming Interface.

ATLAS: Asteroid Terrestrial-impact Last Alert System.

AURA: Association of Universities for Research in Astronomy.

Butler: A middleware component for persisting and retrieving image datasets (raw or processed), calibration reference data, and catalogs.

CADC: Canadian Astronomy Data Centre.

cloud: A visible mass of condensed water vapor floating in the atmosphere, typically high above the ground or in interstellar space acting as the birthplace for stars. Also a way of computing (on other peoples computers leveraging their services and availability).

DAC: Data Access Center.

Data Management System: The computing infrastructure, middleware, and applications that process, store, and enable information extraction from the LSST dataset; the DMS will process peta-scale data volume, convert raw images into a faithful representation of the universe, and archive the results in a useful form. The infrastructure layer consists of the computing, storage, networking hardware, and system software. The middleware layer handles distributed processing, data access, user interface, and system operations services. The applications layer includes the data pipelines and the science data archives' products and services.

DCR: Differential Chromatic Refraction.

DES: Dark Energy Survey.

DOE: Department of Energy.

DP0: Data Preview 0.

DP1: Data Preview 1.

DP2: Data Preview 2.

DR1: Data Release 1.

DR3: Data Release 3.

DRP: Data Release Processing.

EDP2: Early Data Preview 2.

Firefly: A framework of software components written by IPAC for building web-based user interfaces to astronomical archives, through which data may be searched and retrieved, and viewed as FITS images, catalogs, and/or plots. Firefly tools will be integrated into the Science Platform.

HEALPix: Hierarchical Equal-Area iso-Latitude Pixelisation.

HiPS: Hierarchical Progressive Survey (IVOA standard).

IVOA: International Virtual Observatory Alliance.

LSST: Legacy Survey of Space and Time.

LSSTCam: LSST Science Camera.

LSSTComCam: Rubin Commissioning Camera.

metadata: General term for data about data, e.g., attributes of astronomical objects (e.g. images, sources, astroObjects, etc.) that are characteristics of the objects themselves, and facilitate the organization, preservation, and query of data sets (e.g., a FITS header contains metadata).

NSF: National Science Foundation.

Pan-STARRS: Panoramic Survey Telescope and Rapid Response System.

patch: An quadrilateral sub-region of a sky tract, with a size in pixels chosen to fit easily into memory on desktop computers.

POSIX: Portable Operating System Interface.

PPDB: Prompt Products DataBase.

PSF: Point Spread Function.

Qserv: LSST's distributed parallel database. This database system is used for collecting, storing, and serving LSST Data Release Catalogs and Project metadata, and is part of the Software Stack.

REST: REpresentational State Transfer.

RSP: Rubin Science Platform.

S3: (Amazon) Simple Storage Service.

schema: The definition of the metadata and linkages between datasets and metadata entities in a collection of data or archive.

SDSS: Sloan Digital Sky Survey.

shape: In reference to a Source or Object, the shape is a functional characterization of its spatial intensity distribution, and the integral of the shape is the flux. Shape characterizations are a data product in the DIASource, DIAObject, Source, and Object catalogs.

SLR: Single Lens Reflex.

SODA: Server-side Operations for Data Access (IVOA standard).

software: The programs and other operating information used by a computer.

SQL: Structured Query Language.

TAP: Table Access Protocol (IVOA standard).

TOPCAT: Tool for OPerations on Catalogues And Tables.

tract: A portion of sky, a spherical convex polygon, within the LSST all-sky tessellation (sky map). Each tract is subdivided into sky patches.

USDF: United States Data Facility.

VO: Virtual Observatory.

VST: VLT Survey Telescope.

WebDav: Web Distributed Authoring and Versioning.

WFD: Wide-Fast-Deep.

XP: B or R Photometry (Gaia).

DRAFT