



LARGE SYNOPTIC SURVEY TELESCOPE

Large Synoptic Survey Telescope (LSST) Data Products Definition Document

**M. Jurić, T. Axelrod, A.C. Becker, J. Becla, E. Bellm, J.F. Bosch,
D. Ciardi, A.J. Connolly, G.P. Dubois-Felsmann, F. Economou,
M. Freemon, M. Gelman, M. Graham, Ž. Ivezić, T. Jenness,
J. Kantor, K.S. Krughoff, K-T Lim, R.H. Lupton, F. Mueller,
D. Nidever, M. Patterson, D. Petravick, D. Shaw, C. Slater,
M. Strauss, J. Swinbank, J.A. Tyson, M. Wood-Vasey, and X. Wu**

LSE-163

Latest Revision: 2018-01-04

This LSST document has been approved as a Content-Controlled Document. Its contents are subject to configuration control and may not be changed, altered, or their provisions waived without prior approval. If this document is changed or superseded, the new document will retain the Handle designation shown above. The control is on the most recent digital document with this Handle in the LSST digital archive and not printed versions.

Abstract

This document describes the data products and processing services to be delivered by the Large Synoptic Survey Telescope (LSST).

The LSST will deliver three levels of data products and services. **Level 1** (nightly) data products will include images, difference images, catalogs of sources and objects detected in difference images, and catalogs of Solar System objects. Their primary purpose is to enable rapid follow-up of time-domain events. **Level 2** (annual) data products will include well calibrated single-epoch images, deep coadds, and catalogs of objects, sources, and forced sources, enabling static sky and precision time-domain science. **Level 3** (user-created) data product services will enable science cases that greatly benefit from co-location of user processing and/or data within the LSST Archive Center. LSST will also devote 10% of observing time to programs with special cadence. Their data products will be created using the same software and hardware as Levels 1 and 2. All data products will be made available using user-friendly databases and web services.

Change Record

Version	Date	Description	Owner name
1	2013-10-07	Initial release	Mario Juric
2	2016-09-26	Implementation of LCR-758 Update Data Products Definition Document, LSE-163	Gregory Dubois-Felsmann (LCR), Tim Jenness (document), Robert McKercher (Docushare)
3.0	2017-07-03	Implementation of LCR-962. Reference LSE-61 requirements. Minor clean ups.	Tim Jenness
3.1	2018-01-04	Implementation of LCR-1024. Mention relevant OSS requirements.	Tim Jenness

Document curator: Mario Jurić

Document source location: <https://github.com/lstt/LSE-163>

Contents

1	Introduction	1
1.1	The Large Synoptic Survey Telescope	1
1.2	General Image Processing Concepts for LSST	2
1.3	Classes of LSST Data Products	3
2	General Considerations	5
2.1	Estimator and Naming Conventions	6
2.2	Image Characterization Data	7
2.3	Fluxes and Magnitudes	7
2.4	Uniqueness of IDs across database versions	8
2.5	Repeatability of Queries	8
3	Level 1 Data Products	10
3.1	Overview	10
3.2	Level 1 Data Processing	11
3.2.1	Difference Image Analysis	11
3.2.2	Solar System Object Processing	13
3.3	Level 1 Catalogs	14
3.3.1	DIASource Table	15
3.3.2	DIAObject Table	20
3.3.3	SSObject Table	22
3.3.4	Precovery Measurements	23
3.3.5	Reprocessing the Level 1 Data Set	24
3.4	Level 1 Image Products	25
3.4.1	Visit Images	25
3.4.2	Difference Images	25
3.4.3	Image Differencing Templates	25
3.5	Alerts to DIASources	26
3.5.1	Information Contained in Each Alert	26

3.5.2 Receiving and Filtering the Alerts	27
4 Level 2 Data Products	29
4.1 Overview	29
4.2 Level 2 Data Processing	30
4.2.1 Object Characterization Measures	32
4.2.2 Supporting Science Cases Requiring Full Posteriors	34
4.2.3 Source Characterization	35
4.2.4 Forced Photometry	36
4.2.5 Crowded Field Photometry	36
4.3 The Level 2 Catalogs	37
4.3.1 The Object Table	37
4.3.2 Source Table	42
4.3.3 ForcedSource Table	44
4.4 Level 2 Image Products	44
4.4.1 Visit Images	45
4.4.2 Calibration Data	45
4.4.3 Coadded Images	45
4.5 Data Release Availability and Retention Policies	47
5 Level 3 Data Products and Capabilities	48
5.1 Level 3 Data Products and Associated Storage Resources	48
5.2 Level 3 Processing Resources	49
5.3 Level 3 Programming Environment and Framework	50
5.4 Migration of Level 3 data products to Level 2	52
6 Data Products for Special Programs	53
A Appendix: Conceptual Pipeline Design	56

Data Products Definition Document

Preface

The purpose of this document is to describe the data products produced by the Large Synoptic Survey Telescope (LSST).

To a future LSST user, it should clarify what catalogs, image data, software, and services they can expect from LSST. To LSST builders, it provides direction on how to flow down the LSST System Requirements Document to system design, sizing, budget and schedule as they pertain to the data products.

Though under strict change control, this is a ***living document***. LSST will undergo a period of construction and commissioning lasting no less than seven years, followed by a decade of survey operations. To ensure their continued scientific adequacy, the designs and plans for LSST Data Products will be periodically reviewed and updated.

1 Introduction

1.1 The Large Synoptic Survey Telescope

LSST will be a large, wide-field ground-based optical telescope system designed to obtain multiple images covering the sky that is visible from Cerro Pachón in Northern Chile. The current baseline design, with an 8.4m (6.7m effective) primary mirror, a 9.6 deg² field of view, and a 3.2 Gigapixel camera, will allow about 10,000 square degrees of sky to be covered every night using pairs of 15-second exposures, with typical 5σ depth for point sources of $r \sim 24.5$ (AB). The system is designed to yield high image quality as well as superb astrometric and photometric accuracy. The total survey area will include $\sim 30,000$ deg² with $\delta < +34.5^\circ$, and will be imaged multiple times in six bands, *ugrizy*, covering the wavelength range 320–1050 nm. For a more detailed, but still concise, summary of LSST, please see the LSST Overview paper [7]¹.

The project is scheduled to begin the regular survey operations at the start of next decade. About 90% of the observing time will be devoted to a deep-wide-fast survey mode which will uniformly observe a 18,000 deg² region about 1000 times (summed over all six bands) during the anticipated 10 years of operations, and yield a coadded map to $r \sim 27.5$. These data will result in catalogs including over 38 billion stars and galaxies, that will serve the majority of the primary science programs. The remaining 10% of the observing time will be allocated to special projects such as a Very Deep and Fast time domain survey².

The LSST will be operated in fully automated survey mode. The images acquired by the LSST Camera will be processed by LSST Data Management software to a) detect and characterize imaged astrophysical sources and b) detect and characterize temporal changes in the LSST-observed universe. The results of that processing will be reduced images, catalogs of detected objects and the measurements of their properties, and prompt alerts to “events” – changes in astrophysical scenery discovered by differencing incoming images against older, deeper, images of the sky in the same direction (*templates*, see §3.4.3). Measurements will be internally and absolutely calibrated.

The *broad, high-level*, requirements for LSST Data Products are given by the *LSST Science Requirements Document* (SRD; LPM-17). This document lays out the *specifics* of what the data

¹<http://ls.st/2m9>

²Informally known as “Deep Drilling Fields”.

products will comprise of, how those data will be generated, and when. It serves to inform the flow-down from the LSST [SRD](#) through the *LSST System Requirements Document* (the [LSR](#); LSE-29) and the *LSST Observatory System Specifications* ([OSS](#); LSE-30), to the *LSST Data Management System Requirements* ([DMSR](#); LSE-61), the UML model (LDM-133), and the database schema (LDM-153). The [OSS](#) explicitly ties this document to the requirements flow down through requirements OSS-REQ-0126 (Level 1), OSS-REQ-0133 (Level 2), OSS-REQ-0139 (Level 3), OSS-REQ-0391 (general conventions), and OSS-REQ-0392 (special programs). Throughout this document margin notes are used to provide linkage to formal LSST requirements and parameters associated with the nearby text.

1.2 General Image Processing Concepts for LSST

A raw image (baselined as a pair of successive 15-second exposures, called snaps), delivered by the LSST camera, is processed by the Instrument Signature Removal (ISR) pipeline, to produce a single-visit image with, at least conceptually, counts proportional to photon flux entering the telescope pupil (in reality, there are many additional optical, pixel and band-pass effects, including random counting noise and various subtle systematic errors, that are treated during subsequent processing). This single-visit image processed by the ISR is called a “Processed Visit Image” and its main data structures include counts, their variance and various masks, all defined on per pixel basis. After the ISR step is completed, the pixel values and their variance are not modified any more. These single-visit images are used downstream to produce coadded and difference images. The rest of the processing is essentially a model-based interpretation of imaging observations that includes numerous astrophysical and other assumptions.

The basic interpretation model assumes a sum of discrete (but possibly overlapping) sources and a relatively smooth background. The background has a different spectral energy distribution than discrete sources, and it can display both spatial gradients and temporal changes. Discrete sources can vary in brightness and position. The motion can be slow or fast (between two successive observations, less or more motion than about the seeing disk size), and naturally separates stars with proper motions and trigonometric parallax from moving objects in the Solar System. Some objects that vary in brightness can be detectable for only a short period of time (e.g., supernovae and other cosmic explosions).

The image interpretation model separates time-independent model component from a temporally changing component (“DC” and “AC”, respectively). Discrete DC sources are *not* opera-




FIGURE 1: Overview of data products produced by LSST Imaging Processing Science Pipelines.

tionally nor astrophysically associated with discrete AC sources even when they are spatially coincident.

Images (a series of Footprints, where Footprint is a set of connected pixels with counts above some threshold level set by noise properties) of discrete objects are modeled using two models (§4.2.1). A two-component galaxy model includes a linear combination of bulge and disk, with their radial intensity variation described using Sersic profiles. Stars are modeled using a moving point source model with its parallax motion superposed on a linear proper motion. This model shares motion parameters across the six bandpasses and assumes constant flux in each band, and thus includes 11 free parameters. Both galaxy and stellar models are fit to all objects, except for fast-moving objects (the Solar System objects), which are treated separately. Discrete objects detected in *difference* images will be modeled using three models: a point source model, a trailed point source model, and a point source dipole model.

1.3 Classes of LSST Data Products

The main LSST data products are illustrated in Figure 1 (see Appendix for a conceptual design of pipelines which will produce these data products). LSST Data Management will perform two, somewhat overlapping in scientific intent, types of image analyses:

1. Analysis of difference images, with the goal of detecting and characterizing astrophysical phenomena revealed by their time-dependent nature. The detection of supernovae superimposed on bright extended galaxies is an example of this analysis. The processing will be done on nightly or daily basis and result in **Level 1** data products. Level 1 products will include difference images, catalogs of sources detected in difference images (DIASources), astrophysical objects³ these are associated to (DIAObjects), and Solar System objects (SSObjects⁴). The catalogs will be entered into the **Level 1 database** and made available in near real time. Notifications (“alerts”) about new DIASources will be issued using community-accepted standards⁵ within 60 seconds of observation. Level 1 data products are discussed in § 3.
2. Analysis of direct images, with the goal of detecting and characterizing astrophysical objects. Detection of faint galaxies on deep coadds and their subsequent characterization is an example of this analysis. The results are **Level 2** data products. These products, generated and released annually⁶, will include the single-epoch images, deep coadds, catalogs of characterized Objects (detected on deep coadds as well as individual visits⁷), Sources⁸ (detections and measurements on individual visits), and ForcedSources (constrained measurement of flux on individual visits). It will also include fully reprocessed Level 1 data products (see §3.3.5). In contrast to the “living” Level 1 database, which is updated in real-time, the Level 2 databases are static and will not change after release. Level 2 data products are discussed in § 4.

The two types of analyses have different requirements on timeliness. Changes in flux or position of objects may need to be immediately followed up, lest interesting information be lost. Thus the primary results of analysis of difference images – discovered and characterized DIASources – generally need to be broadcast as *event alerts* within 60 seconds of end of visit acquisition. The analysis of science (direct) images is less time sensitive, and will be done as

OTT1

DMS-REQ-0004

³The LSST has adopted the nomenclature by which single-epoch detections of astrophysical *objects* are called *sources*. The reader is cautioned that this nomenclature is not universal: some surveys call *detections* what LSST calls *sources*, and use the term *sources* for what LSST calls *objects*.

⁴SSObjects used to be called “Moving Objects” in previous versions of the LSST Data Products baseline. The name is potentially confusing as high-proper motion stars are moving objects as well. A more accurate distinction is the one between objects *inside* and *outside* of the Solar System.

⁵For example, VOEvent, see <http://ls.st/4tt>

⁶Except for the first two data releases, which will be created six months apart.

⁷The LSST takes two exposures per pointing, nominally 15 seconds in duration each, called *snaps*. For the purpose of data processing, that pair of exposures will typically be coadded and treated as a single exposure, called a *visit*.

⁸When written in bold monospace type (i.e., \tt), Objects and Sources refer to objects and sources detected and measured as a part of Level 2 processing.

a part of annual data release process.

Recognizing the diversity of astronomical community needs, and the need for specialized processing not part of the automatically generated Level 1 and 2 products, LSST plans to devote 10% of its data management system capabilities to enabling the creation, use, and federation of **Level 3** (user-created) data products. Level 3 capabilities will enable science cases that greatly benefit from co-location of user processing and/or data within the LSST Archive Center. The high-level requirement for Level 3 is established in § 3.5 of the LSST [SRD](#). Their details are discussed in § 5 of this document.

Finally, LSST Survey Specifications (§ 3.4 of LSST [SRD](#)) prescribe that 90% of LSST observing time be spent in the so-called “universal cadence” mode of surveying the sky. These observations will result in Level 1 and 2 data products discussed above. The remaining 10% of observing time will be devoted to **special programs**, designed to obtain improved coverage of interesting regions of observational parameter space. Examples include very deep ($r \sim 26$, per exposure) observations, observations with very short revisit times (~ 1 minute), and observations of “special” regions such as the Ecliptic, Galactic plane, and the Large and Small Magellanic Clouds. The data products for these programs will be generated using the same processing software and hardware and possess the general characteristics of Level 1 and 2 data products, but may be performed on a somewhat different cadence. They will be discussed in § 6.

2 General Considerations

Most LSST data products will consist of images and/or catalogs. The catalogs will be stored and offered to the users as *relational databases* which they will be able to query. This approach was shown to work well by prior large surveys, for example the Sloan Digital Sky Survey (SDSS).

Different data products will generally be stored in different databases. For example, Level 1 data products will be stored in a *Level 1 database*. Level 2 “universal cadence” products will be stored in a *Level 2 database*. The products for special programs may be stored in many different databases, depending on the nature of the program.

Nevertheless, all these databases will follow certain naming and other conventions. We discuss these in the subsections to follow.

2.1 Estimator and Naming Conventions

For all catalogs data, we will employ a convention where estimates of standard errors have the suffix `Err`, while the estimates of inherent widths of distribution (or functions in general) have the suffix `Sigma`⁹. The latter are defined as the square roots of the second moment about the quoted value of the quantity at hand.

DMS-REQ-0333

Unless noted otherwise, maximum likelihood values (called likelihood for simplicity) will be quoted for all fitted parameters (measurements). Together with covariances, these let the end-user apply whatever prior they deem appropriate when computing posteriors¹⁰. Where appropriate, multiple independent samples from the likelihood may be provided to characterize departures from Gaussianity.

We will provide values of log likelihood, the χ^2 for the fitted parameters, and the number of data points used in the fit. Database functions (or precomputed columns) will be provided for frequently used combinations of these quantities (e.g., χ^2/dof). These can be used to assess the model fit quality. Note that, *if the errors of measured quantities are normally distributed*, the likelihood is related to the χ^2 as:

DMS-REQ-0331

$$L = \left(\prod_k \frac{1}{\sigma_k \sqrt{2\pi}} \right) \exp \left[-\frac{\chi^2}{2} \right] \quad (1)$$

where the index k runs over all data points included in the fit. For completeness, χ^2 is defined as:

$$\chi^2 = \sum_k \left(\frac{x_k - \bar{x}}{\sigma_k} \right)^2, \quad (2)$$

where $\bar{x}$ is the mean value of x_k .

For fluxes, we recognize that a substantial fraction of astronomers will just want the posteriors marginalized over all other parameters, trusting the LSST experts to select an appropriate prior¹¹. For example, this is nearly always the case when constructing color-color or color-magnitude diagrams. We will support these use cases by providing additional pre-computed columns, taking care to name them appropriately so as to minimize accidental incorrect us-

DMS-REQ-0331

⁹Given N measurements, standard errors scale as $N^{-1/2}$, while widths remain constant.

¹⁰There's a tacit assumption that a Gaussian is a reasonably good description of the likelihood surface around the ML peak.

¹¹It's likely that most cases will require just the expectation value alone.

age. For example, a column named `gFlux` may be the expectation value of the g-band flux, while `gFluxML` may represent the maximum likelihood value.

2.2 Image Characterization Data

Raw images will be processed to remove instrumental signature and characterize their properties, including backgrounds (both due to night sky and astrophysical), the point spread function and its variation, photometric zero-point model, and the world coordinate system (WCS).

DMS-REQ-0069

DMS-REQ-0327

DMS-REQ-0070

DMS-REQ-0029

DMS-REQ-0030

That characterization is crucial for deriving LSST catalogs and understanding the images. It will be kept and made available to the users. The exact format used to store these (meta)data will depend on the final adopted algorithm in consultation with the scientific community to ensure the formats in which these data are served are maximally useful.

DMS-REQ-0328

DMS-REQ-0072

Each processed image¹², including the coadds, will record information on pixel variance (the “variance plane”), as well as per-pixel masks (the “mask plane”). These will allow the users to determine the validity and usefulness of each pixel in estimating the flux density recorded in that area of the sky.

This information will be per-pixel, and potentially unwieldy to use for certain science cases. We plan to investigate approximate schemes for storing masks based on geometry (e.g., similar to Mangle or STOMP), *in addition* to storing them on a per pixel basis.

DMS-REQ-0326

2.3 Fluxes and Magnitudes

DMS-REQ-0043

Because flux measurements on difference images (Level 1 data products; § 3) are performed against a template and thus represent a flux difference, the measured flux of a source on the difference image can be negative. The flux can also go negative for faint sources in the presence of noise. Negative fluxes cannot be stored as (Pogson) magnitudes; $\log$ of a negative number is undefined. We therefore prefer to store fluxes, rather than magnitudes, in database tables¹³.

DMS-REQ-0347

We quote fluxes in units of “maggie”. A maggie, as introduced by SDSS, is a linear measure of

¹²It is also frequently referred to as *calibrated exposure*, from the Butler product type of `calexp`.

¹³This is a good idea in general. E.g., given multi-epoch observations, one should always be averaging fluxes, rather than magnitudes.

flux. It is defined so that an object having a flux of one maggie (integrated over the bandpass) has an AB magnitude of zero:

$$m_{AB} = -2.5 \log_{10}(f/\text{maggie}) \quad (3)$$

We chose to use maggies (as opposed to, say, Jansky) to allow the user to differentiate between two distinct sources of photometric calibration error: the error in relative (internal) calibration of the survey, and the error in absolute calibration that depends on the knowledge of absolute flux of photometric standards.

Nevertheless, we acknowledge that the large majority of users will want to work with magnitudes. For convenience, we plan to provide columns with (Pogson) magnitudes¹⁴, where values with negative flux will evaluate to NULL. Similarly, we will provide columns with flux expressed in Jy (and its error estimate that includes the relative and absolute calibration error contributions).

2.4 Uniqueness of IDs across database versions

DMS-REQ-0292

To reduce the likelihood for confusion, all IDs shall be unique across databases and database versions, other than those corresponding to uniquely identifiable entities (i.e., IDs of exposures).

For example, DR4 and DR5 (or any other) release will share no identical Object, Source, DIA-Object or DIASource IDs (see § 3 and 4 for the definitions of Objects, DIAObjects, etc.).

2.5 Repeatability of Queries

DMS-REQ-0291

We require that queries executed at a known point in time against any LSST-delivered database be repeatable at a later date. This promotes the reproducibility of science derived from LSST data. It is of special importance for Level 1 catalogs (§ 3) that will change on a nightly basis as new time domain data is being processed and added to the catalogs.

The exact implementation of this requirement is left to the LSST Data Management database

¹⁴These will most likely be implemented as “virtual” or “computed” columns

team. One possibility may be to make the key tables (nearly) append-only, with each row having two timestamps – createdTai and deletedTai, so that queries may be limited by a WHERE clause:

```
SELECT * FROM DIASource WHERE 'YYYY-MM-DD-HH-mm-SS '  
      BETWEEN createdTAI and deletedTAI
```

or, more generally:

```
SELECT * FROM DIASource WHERE "data is valid as of YYYY-MM-DD"
```

A perhaps less error-prone alternative, if technically feasible, may be to provide multiple virtual databases that the user would access as:

```
CONNECT lsst-dr5-yyyy-mm-dd  
SELECT * FROM DIASource
```

The latter method would probably be limited to nightly granularity, unless there's a mechanism to create virtual databases/views on-demand.

3 Level 1 Data Products

3.1 Overview

Level 1 data products are a result of difference image analysis (DIA; §3.2.1). They include the sources detected in difference images (DIASources), astrophysical objects that these are associated to (DIAObjects), identified Solar System objects¹⁵ (SSObject), and related, broadly defined, metadata (including eg., cut-outs¹⁶).

DIASources are sources detected on difference images (those with the signal-to-noise ratio $S/N > transSNR$ after correlation with the PSF profile, with $transSNR$ defined in the SRD and presently set to 5). They represent changes in flux with respect to a deep template. Detections with high probability of being instrumental non-astrophysical artifacts may be excluded. Physically, a DIASource may be an observation of new astrophysical object that was not present at that position in the template image (for example, an asteroid), or an observation of flux change in an existing source (for example, a variable star). Their flux can be negative (eg., if a source present in the template image reduced its brightness, or moved away). Their shape can be complex (eg., trailed, for a source with proper motion approaching $\sim \text{deg/day}$, or “dipole-like”, if an object’s observed position exhibits an offset – true or apparent – compared to its position on the template). Some DIASources will be caused by background fluctuations; with $transSNR = 5$, we expect about one such false positive per CCD (of the order 200,000 per typical night). The expected number of false positives due to background fluctuations is a very strong function of adopted $transSNR$: a change of $transSNR$ by 0.5 results in a variation of an order of magnitude, and a change of $transSNR$ by unity changes the number of false positives by about two orders of magnitude.

$transSNR$

Clusters of DIASources detected on visits taken at different times are associated with either a DIAObject or an SSObject, to represent the underlying astrophysical phenomenon. The association can be made in two different ways: by assuming the underlying phenomenon is an object within the Solar System moving on an orbit around the Sun¹⁷, or by assuming it

DMS-REQ-0285

¹⁵The SRD considers Solar System object orbit catalog to be a Level 2 data product (LSST SRD, Sec 3.5). Nevertheless, to successfully differentiate between apparitions of known Solar System objects and other types DIASources we consider it functionally a part of Level 1.

¹⁶Small, 30×30 , sub-images at the position of a detected source. Also known as *postage stamps*.

¹⁷We don’t plan to fit for motion around other Solar System bodies; eg., identifying new satellites of Jupiter is left to the community.

to be distant enough to only exhibit small parallactic and proper motion¹⁸. The latter type of association is performed during difference image analysis right after the image has been acquired. The former is done at daytime by the Moving Objects Processing Software (MOPS), unless the DIASource is an apparition of an already known SSObject. In that case, it will be flagged as such during difference image analysis.

DMS-REQ-0089

DMS-REQ-0002

At the end of the difference image analysis of each visit, we will issue time domain event alerts for all newly detected DIASources¹⁹.

3.2 Level 1 Data Processing

3.2.1 Difference Image Analysis

The following is a high-level description of steps which will occur during regular *nightly* difference image analysis (see Figures 3 and 5):

1. A visit is acquired and reduced to a single *Processed Visit Image* (cosmic ray rejection, instrumental signature removal²⁰, combining of snaps, etc.).
2. The Processed Visit Image is differenced against the appropriate template and DIASources are detected. If necessary, deblending will be performed at this stage. Both the parent blend and the deblended children will be measured and stored as DIASources (see next item), but only the children will be matched against DIAObjects and alerted on. Deblended objects will be flagged as such.
3. The flux and shape²¹ of the DIASource are measured on the difference image. PSF photometry is performed on the Processed Visit Image at the position of the DIASource to obtain a measure of the total flux.
4. The Level 1 database (see §3.3) is searched for a DIAObject or an SSObject with which

DMS-REQ-0069

DMS-REQ-0010

DMS-REQ-0269

¹⁸Where 'small' is small enough to unambiguously positionally associate together individual apparitions of the object.

¹⁹For observations on the Ecliptic near the opposition Solar System objects will dominate the DIASource counts and (until they're recognized as such) overwhelm the explosive transient signal. It will therefore be advantageous to quickly identify the majority of Solar System objects early in the survey.

²⁰Eg., subtraction of bias and dark frames, flat fielding, bad pixel/column interpolation, etc.

²¹The "shape" in this context consists of weighted 2nd moments, as well as fits to a trailed source model and a dipole model.

to positionally associate the newly discovered DIASource²². If no match is found, a new DIAObject is created and the observed DIASource is associated to it.

DMS-REQ-0273

DMS-REQ-0271

DMS-REQ-0285

5. If the DIASource has been associated with an SSObject (a known Solar System object), it will be flagged as such and an alert will be issued. Further processing will occur in daytime (see section 3.2.2).

DMS-REQ-0274

6. Otherwise, the associated DIAObject measurements will be updated with new data collected during previous 12 months. Hence, the computed parameters for DIAObjects have a “memory” of past data that does not extend beyond this cutoff²³. All affected columns will be recomputed, including proper motions, centroids, light curves, etc.

DMS-REQ-0319

diaCharacterization-
Cutoff

7. The Level 2 database²⁴ is searched for the three nearest stars and three nearest galaxies to the DIAObject in Objects out to some maximum radius. The IDs of these nearest-neighbor Objects are recorded in the DIAObject record and provided in the issued event alert (see below).

diaNearbyObjMaxS-
tar
diaNearbyObjMax-
Galaxy
diaNearbyObjRadius

DMS-REQ-0271

DMS-REQ-0002

8. An alert is issued that includes: the name of the Level 1 database, the timestamp of when this database has been queried to issue this alert, the DIASource ID, the SSObject ID or DIAObject ID²⁵, name of the Level 2 database and the IDs of nearby Objects, and the associated science content (centroid, fluxes, low-order lightcurve moments, periods, etc.), *including all* DIASources from the last 12 months that are linked with the SSObject or DIAObject. The science content associated with the Level 2 database Objects will not be included. See Section 3.5 for a more complete enumeration.

diaCharacterization-
Cutoff

DMS-REQ-0274

9. For all DIAObjects overlapping the field of view, including those that have an associated new DIASource from this visit, forced photometry will be performed on difference image (point source photometry only). Those measurements will be stored as appropriately flagged DIASources²⁶. No alerts will be issued for these DIASources.

DMS-REQ-0317

10. Within 24 hours of discovery, *precovery* PSF forced photometry will be performed on

L1PublicT

²²The association algorithm will guarantee that a DIASource is associated with not more than one existing DIAObject or SSObject. The algorithm will take into account the parallax and proper (or Keplerian) motions, as well as the errors in estimated positions of DIAObject, SSObject, and DIASource, to find the maximally likely match. Multiple DIASources in the same visit will not be matched to the same DIAObject.

²³This restriction is removed when Level 1 processing is rerun during Data Release production, see § 3.3.5.

²⁴Level 2 database is a database resulting from annual data release processing. See § 4 for details.

²⁵We guarantee that a receiver will always be able to regenerate the alert contents at any later date using the included timestamps and metadata (IDs and database names).

²⁶For the purposes of this document, we’re treating the DIASources generated by forced photometry or precovery measurements to be the same as DIASources detected in difference images (but flagged appropriately). In the logical schema, these may be divided into two separate tables.

any difference image overlapping the position of new DIAObjects taken within the past 30 days, and added to the database. Alerts will not be issued with precovery photometry information.

precoveryWindow

DMS-REQ-0287

In addition to the processing described above, a smaller sample of sources detected on difference images *below* the nominal $transSNR = 5$ threshold will be measured and stored, in order to enable monitoring of difference image analysis quality.

transSNR

DMS-REQ-0270

Also, the system will have the ability to measure and alert on a limited²⁷ number of sources detected below the nominal threshold for which additional criteria are satisfied. For example, a $transSNR = 3$ source detection near a gravitational keyhole²⁸ may be highly significant in assessing the danger posed by a potentially hazardous asteroid. The initial set of criteria will be defined by the start of LSST operations.

3.2.2 Solar System Object Processing

The following will occur during regular Solar System object processing (in daytime²⁹, after a night of observing; see Figure 6) :

DMS-REQ-0004

DMS-REQ-0089

1. The orbits and physical properties of all SSObjects re-observed on the previous night are recomputed. External orbit catalogs (or observations) are also used to improve orbit estimates. Updated data are entered to the SSObjects table.
2. All DIASources detected on the previous night, that have *not* been matched at a high confidence level to a known Object, DIAObject, SSObject, or an artifact, are analyzed for potential pairs, forming *tracklets*.
3. The collection of tracklets collected over the past 30 days³⁰ is searched for subsets forming *tracks* consistent with being on the same Keplerian orbit around the Sun.

DMS-REQ-0288

DMS-REQ-0273

²⁷It will be sized for no less than ~10% of average DIASource per visit rate.

²⁸A gravitational keyhole is a region of space where Earth's gravity would modify the orbit of a passing asteroid such that the asteroid would collide with the Earth in the future.

²⁹Note that there is *no strict bound on when daytime Solar System processing must finish*, just that, averaged over some reasonable timescale (eg., a month), a night's worth of observing is processed within 24 hours. Nights rich in moving objects may take longer to process, while nights with less will finish more quickly. In other words, the requirement is on *throughput*, not latency.

³⁰The exact time period is a design and computational question discussed in LDM-156.

4. For those that are, an orbit is fitted and a new SSObject table entry created. DIASource records are updated to point to the new SSObject record. DIAObjects “orphaned” by this unlinking are deleted.³¹.
5. Preccovery linking is attempted for all SSObjects whose orbits were updated in this process. Where successful, SSObjects (orbits) are recomputed as needed.

DMS-REQ-0286

3.3 Level 1 Catalogs

The described alert processing design relies on the “living” Level 1 database that contains the objects and sources detected on difference images. At the very least³², this database will have tables of DIASources, DIAObjects, and SSObjects, populated in the course of nightly and daily difference image and Solar System object processing³³. As these get updated and added to, their updated contents becomes visible (query-able) immediately³⁴.

DMS-REQ-0269

DMS-REQ-0271

DMS-REQ-0273

DMS-REQ-0312

This database is *only loosely coupled to the Level 2 database*. All of the coupling is through positional matches between the DIAObjects entries in the Level 1 database and the Objects in the Level 2 database database. There is no direct DIASource-to-Object match: in general, a time-domain object is not necessarily the same astrophysical object as a static-sky object, even if the two are positionally coincident (eg. an asteroid overlapping a galaxy). Therefore, adopted data model emphasizes that *having a DIASource be positionally coincident with an Object does not imply it is physically related to it*. Absent other information, the least presumptuous data model relationship is one of *positional association*, not *physical identity*.

This may seem odd at first: for example, in a simple case of a variable star, matching individual DIASources to Objects is exactly what an astronomer would want. That approach, however, fails in the following scenarios:

- *A supernova in a galaxy*. The matched object in the Object table will be the galaxy, which is a distinct astrophysical object. We want to keep the information related to the supernova (eg., colors, the light curve) separate from those measurements for the galaxy.

³¹Some DIAObjects may only be left with forced photometry measurements at their location (since all DIAObjects are force-photometered on previous and subsequent visits); these will be kept but flagged as such.

³²It will also contain exposure and visit metadata, MOPS-specific tables, etc. These are either standard/uncontroversial, implementation-dependent, or less directly relevant for science and therefore not discussed in this document.

³³The latter is also colloquially known as *DayMOPS*.

³⁴No later than the moment of issuance of any event alert that may refer to it.

- *An asteroid occulting a star.* If associated with the star on first apparition, the association would need to be dissolved when the source is recognized as an asteroid (perhaps even as early as a day later).
- *A supernova on top of a pair of blended galaxies.* It is not clear in general to which galaxy this DIASource would “belong”. That in itself is a research question.

DIASource-to-Object matches can still be emulated via a two-step relation (DIASource-DIA-Object-Object). For ease of use, views or pre-built table with these matches will be offered to the end-users.

DMS-REQ-0324

In the sections to follow, we present the *conceptual schemas* for the most important Level 1 database tables. These convey *what* data will be recorded in each table, rather than the details of *how*. For example, columns whose type is an array (eg., `radec`) may be expanded to one table column per element of the array (eg., `ra`, `dec1`) once this schema is translated to SQL³⁵. Secondly, the tables to be presented are largely normalized (i.e., contain no redundant information). For example, since the band of observation can be found by joining a DIA-Source table to the table with exposure metadata, there’s no column named `band` in the DIA-Source table. In the as-built database, the views presented to the users will be appropriately denormalized for ease of use.

3.3.1 DIASource Table

This is a table of sources detected at $transSNR \geq 5$ on difference images³⁶ (DIASources). On average, the LSST SRD expects up to $\sim 10,000$ astrophysical DIASources per visit ($\sim 10M$ per night; $100,000$ per deg^2 of the sky per hour).

transSNR
LSR-REQ-0101

transN

Some $transSNR \geq 5$ sources will not be caused by observed astrophysical phenomena, but by artifacts (bad columns, diffraction spikes, etc.). The difference image analysis software will attempt to identify and flag these as such.

Unless noted otherwise, all DIASource quantities (fluxes, centroids, etc.) are measured on the difference image.

DMS-REQ-0269

³⁵The SQL realization of this schema is defined in the `cat` package and documented in LDM-153 and can be browsed at <http://ls.st/8g4>

³⁶This requirement is specified in the LSST SRD.

Table 1: DIASource Table

Name	Type	Unit	Description
diaSourceId	uint64		Unique source identifier
ccdVisitId	uint64		ID of CCD and visit where this source was measured
diaObjectId	uint64		ID of the DIAObject this source was associated with, if any.
ssObjectId	uint64		ID of the SSObject this source has been linked to, if any.
parentDiaSourceId	uint64		ID of the parent DIASource this object has been deblended from, if any.
midPointTai	double	time	Time of mid-exposure for this DIA-Source ³⁷ .
radec	double[2]	degrees	Centroid, (α, δ) ³⁸ .
radecCov	float[3]	various	radec covariance matrix.
xy	float[2]	pixels	Column and row of the centroid.
xyCov	float[3]	various	Centroid covariance matrix.
apFlux	float	nmgy	Calibrated aperture flux. Note that this actually measures the flux <i>difference</i> between the template and the visit image.
apFluxErr	float	nmgy	Estimated uncertainty of apFlux.
SNR	float		The signal-to-noise ratio at which this source was detected in the difference image. ³⁹
psFlux	float	nmgy ⁴⁰	Calibrated flux for point source model. Note this actually measures the flux <i>difference</i> between the template and the visit image.

Continued on next page

³⁷The visit mid-exposure time generally depends on the position of the source relative to the shutter blade motion trajectory.

³⁸The astrometric reference frame will be chosen closer to start of operations.

³⁹This is not necessarily the same as apFlux/apFluxErr, as the flux measurement algorithm may be more accurate than the detection algorithm.

⁴⁰A “maggie”, as introduced by SDSS, is a linear measure of flux in units of 3631 Jy; one maggie has an AB magnitude of 0, $m_{AB} = -2.5 \log_{10}(\text{maggie})$. “nmgy” is short for a nanomaggie (1 nmgy = 3.631 μ Jy). For example,

Table 1: DIASource Table

Name	Type	Unit	Description
psRadec	double[2]	degrees	Centroid for point source model.
psCov	float[6]	various	Covariance matrix for point source model parameters.
psLnL	float		Natural <i>log</i> likelihood of the observed data given the point source model.
psChi2	float		χ^2 statistic of the model fit.
psNdata	int		The number of data points (pixels) used to fit the model.
trailFlux	float	nmgy	Calibrated flux for a trailed source model ^{41,42} . Note this actually measures the flux <i>difference</i> between the template and the visit image.
trailRadec	double[2]	degrees	Centroid for trailed source model.
trailLength	float	arcsec	Maximum likelihood fit of trail length ^{43,44} .
trailAngle	float	degrees	Maximum likelihood fit of the angle between the meridian through the centroid and the trail direction (bearing, direction of motion).
trailCov	float[15]	various	Covariance matrix of trailed source model parameters.
trailLnL	float		Natural <i>log</i> likelihood of the observed data given the trailed source model.
trailChi2	float		χ^2 statistic of the model fit.

Continued on next page

a flux of 0.158 nmgy corresponds to AB magnitude of 24.5. See §2.3 for details.

⁴¹A *Trailed Source Model* attempts to fit a (PSF-convolved) model of a point source that was trailed by a certain amount in some direction (taking into account the two-snap nature of the visit, which may lead to a dip in flux around the mid-point of the trail). Roughly, it's a fit to a PSF-convolved line. The primary use case is to characterize fast-moving Solar System objects.

⁴²This model does not fit for the *direction* of motion; to recover it, we would need to fit the model to separately to individual snaps of a visit. This adds to system complexity, and is not clearly justified by increased MOPS performance given the added information.

⁴³Note that we'll likely measure trailRow and trailCol, and transform to trailLength/trailAngle (or trailRa/trailDec) for storage in the database. A stretch goal is to retain both.

⁴⁴TBD: Do we need a separate trailCentroid? It's unlikely that we do, but one may wish to prove it.

Table 1: DIASource Table

Name	Type	Unit	Description
trailNdata	int		The number of data points (pixels) used to fit the model.
dipMeanFlux	float	nmgy	Maximum likelihood value for the mean absolute flux of the two lobes for a dipole model ⁴⁵ .
dipFluxDiff	float	nmgy	Maximum likelihood value for the difference of absolute fluxes of the two lobes for a dipole model.
dipRadec	double[2]	degrees	Centroid for dipole model.
dipLength	float	arcsec	Maximum likelihood value for the lobe separation in dipole model.
dipAngle	float	degrees	Maximum likelihood fit of the angle between the meridian through the centroid and the dipole direction (bearing, from negative to positive lobe).
dipCov	float[21]	various	Covariance matrix of dipole model parameters.
dipLnL	float		Natural <i>log</i> likelihood of the observed data given the dipole source model.
dipChi2	float		χ^2 statistic of the model fit.
dipNdata	int		The number of data points (pixels) used to fit the model.
totFlux	float	nmgy	Calibrated flux for point source model measured on the visit image centered at the centroid measured on the difference image (forced photometry flux)
totFluxErr	float	nmgy	Estimated uncertainty of fpFlux.

Continued on next page

⁴⁵A *Dipole Model* attempts to fit a (PSF-convolved) model of two point sources, with fluxes of opposite signs, separated by a certain amount in some direction. The primary use case is to characterize moving stars and problems with image differencing (e.g., due to astrometric offsets).

Table 1: DIASource Table

Name	Type	Unit	Description
diffFlux	float	nmgy	Calibrated flux for point source model centered on radec but measured on the difference of snaps comprising this visit ⁴⁶ .
diffFluxErr	float	nmgy	Estimated uncertainty of diffFlux.
fpBkgd	float	nmgy/asec ²	Estimated background at the position (centroid) of the object in the template image.
fpBkgdErr	float	nmgy/asec ²	Estimated uncertainty of fpBkgd.
lxx	float	asec ²	Adaptive second moment of the source intensity. See Bernstein & Jarvis [2] for detailed discussion of all adaptive-moment related quantities ⁴⁷ .
lyy	float	asec ²	Adaptive second moment of the source intensity.
lxy	float	asec ²	Adaptive second moment of the source intensity.
lcov	float[6]	asec ⁴	Ixx, Iyy, Ixy covariance matrix.
lxxPSF	float	asec ²	Adaptive second moment for the PSF.
lyyPSF	float	asec ²	Adaptive second moment for the PSF.
lxyPSF	float	asec ²	Adaptive second moment for the PSF.

Continued on next page

⁴⁶This flux can be used to detect sources changing on timescales comparable to snap exposure length (~15 sec).

⁴⁷Or <http://ls.st/5f4> for a brief summary.

Table 1: DIASource Table

Name	Type	Unit	Description
extendedness	float		A measure of extendedness, computed using a combination of available moments, or from a likelihood ratio of point/trailed source models (exact algorithm TBD). <i>extendedness</i> = 1 implies a high degree of confidence that the source is extended. <i>extendedness</i> = 0 implies a high degree of confidence that the source is point-like.
spuriousness	float		A measure of spuriousness, computed using information ⁴⁸ from the source and image characterization, as well as the information on the Telescope and Camera system (e.g., ghost maps, defect maps, etc.).
flags	bit[64]	bit	Various useful flags.

Some fast-moving, trailed, sources may be due to passages of nearby asteroids. Their trails may exhibit significant curvature. While we do not measure the curvature directly, it can be inferred by examining the length of the trail, the trailed model covariance matrices, and the adaptive shape measures. Once curvature is suspected, the users may fit curved trail models to the cutout provided with the alert.

3.3.2 DIAObject Table

DMS-REQ-0271
DMS-REQ-0272
diaNearbyMaxObj

⁴⁸The computation of spuriousness will be “prior free” to the extent possible and not use any information about the astrophysical neighborhood of the source, whether it has been previously observed or not, etc. The intent is to avoid introducing a bias against unusual sources or sources discovered in unusual environments.

Table 2: DIAObject Table

Name	Type	Unit	Description
diaObjectId	uint64		Unique identifier.
radec	double[2]	degrees	(α, δ) position of the object at time radecTai.
radecCov	float[3]	various	radec covariance matrix.
radecTai	double	time	Time at which the object was at a position radec.
pm	float[2]	mas/yr	Proper motion vector ⁴⁹ .
parallax	float	mas	Trigonometric arallax.
pmParallaxCov	float[6]	various	Proper motion - parallax covariances.
pmParallaxLnL	float		Natural log of the likelihood of the linear proper motion-parallax fit ⁵⁰ .
pmParallaxChi2	float		χ^2 statistic of the model fit.
pmParallaxNdata	int		The number of data points used to fit the model.
psFluxMean	float[ugrizy]	nmgy	Weighted mean of point-source model flux, psFlux.
psFluxMeanErr	float[ugrizy]	nmgy	Standard error of psFluxMean.
psFluxSigma	float[ugrizy]	nmgy	Standard deviation of the distribution of psFlux.
psFluxChi2	float[ugrizy]		χ^2 statistic for the scatter of psFlux around psFluxMean.
psFluxNdata	int[ugrizy]		The number of data points used to compute psFluxChi2.
fpFluxMean	float[ugrizy]	nmgy	Weighted mean of forced photometry flux, fpFlux.
fpFluxMeanErr	float[ugrizy]	nmgy	Standard error of fpFluxMean.
fpFluxSigma	float[ugrizy]	nmgy	Standard deviation of the distribution of fpFlux.

Continued on next page

⁴⁹High proper-motion or parallax objects will appear as “dipoles” in difference images. Great care will have to be taken not to misidentify these as subtraction artifacts.

⁵⁰radec, pm, and parallax will all be simultaneously fitted for.

Table 2: DIAObject Table

Name	Type	Unit	Description
lcPeriodic	float[6 × 32]		Periodic features extracted from light-curves using generalized Lomb-Scargle periodogram [Table 4, 12] ⁵¹ .
lcNonPeriodic	float[6 × 20]		Non-periodic features extracted from light-curves [Table 5, 12].
nearbyObj	uint64[6]		Closest Objects (3 stars and 3 galaxies) in Level 2 database.
nearbyObjDist	float[6]	arcsec	Distances to nearbyObj.
nearbyObjLnP	float[6]		Natural log of the probability that the observed DIAObject is the same as the nearby Object ⁵² .
flags	bit[64]	bit	Various useful flags.

3.3.3 SSObject Table

DMS-REQ-0273

Table 3: SSObject Table

Name	Type	Unit	Description
ssObjectId	uint64		Unique identifier.
oe	double[7]	various	Osculating orbital elements at epoch ($q, e, i, \Omega, \omega, M_0$, epoch).
oeCov	double[28]	various	Covariance matrix for oe.
arc	float	days	Arc of observation.
orbFitLnL	float		Natural log of the likelihood of the orbital elements fit.
orbFitChi2	float		χ^2 statistic of the orbital elements fit.
orbFitNdata	int		The number of data points (observations) used to fit the orbital elements.

Continued on next page

⁵¹The exact features in use when LSST begins operations are likely to be different compared to the baseline described here. This is to be expected given the rapid pace of research in time domain astronomy. However, the *number* of computed features is unlikely to grow beyond the present estimate.

⁵²This quantity will be computed by marginalizing over the product of position and proper motion error ellipses of the Object and DIAObject, assuming an appropriate prior.

Table 3: SSObject Table

Name	Type	Unit	Description
MOID	float[2]	AU	Minimum orbit intersection distances ⁵³
moidLon	double[2]	degrees	MOID longitudes.
H	float[6]	mag	Mean absolute magnitude, per band [10, magnitude-phase system].
G_1	float[6]	mag	G_1 slope parameter, per band [10, magnitude-phase system].
G_2	float[6]	mag	G_2 slope parameter, per band [10, magnitude-phase system].
hErr	float[6]	mag	Uncertainty of H estimate.
g1Err	float[6]	mag	Uncertainty of G_1 estimate.
g2Err	float[6]	mag	Uncertainty of G_2 estimate.
flags	bit[64]	bit	Various useful flags.

The G_1 and G_2 parameters for the large majority of asteroids will not be well constrained until later in the survey. We may decide not to fit for it at all over the first few DRs and add it later in Operations, or provide two-parameter G_{12} fits. Alternatively, we may fit it using strong priors on slopes poorly constrained by the data. The design of the data management system is insensitive to this decision, making it possible to postpone it to Commissioning to ensure it follows the standard community practice at that time.

The LSST database will provide functions to compute the phase (Sun-Asteroid-Earth) angle α for every observation, as well as the reduced, $H(\alpha)$, and absolute, H , asteroid magnitudes in LSST bands.

DMS-REQ-0323

3.3.4 Preccovery Measurements

When a new DIASource is detected, it's useful to perform PSF photometry at the location of the new source on images taken prior to discovery. These are colloquially known as *precovery*

⁵³<http://www2.lowell.edu/users/elgb/moid.html>

*measurements*⁵⁴. Performing precovery in real time over all previously acquired visits is too I/O intensive to be feasible. We therefore plan the following:

DMS-REQ-0287
DMS-REQ-0286

1. For all newly discovered objects, perform precovery PSF photometry on visits taken over the previous 30 days⁵⁵.
2. Make available a “precovery service” to request precovery for a limited number of DIA-Sources across all previous visits, and make it available within 24 hours of the request. Web interface and machine-accessible APIs will be provided.

precoveryWindow

DMS-REQ-0341

The former should satisfy the most common use cases (eg., SNe), while the latter will provide an opportunity for more extensive yet timely precovery of targets of special interest.

3.3.5 Reprocessing the Level 1 Data Set

In what we’ve described so far, the “living” Level 1 database is continually being added to as new images are taken and DIASources identified. Every time a new DIASource is associated to an existing DIAObject, the DIAObject record is updated to incorporate new information brought in by the DIASource. Once discovered and measured, the DIASources would never be re-discovered and re-measured at the pixel level.

DMS-REQ-0312

This would be far from optimal. The instrument will be better understood with time. Newer versions of LSST pipelines will improve detection and measurements on older data. Also, precovery photometry should optimally be performed for *all* objects, and not just a select few. This argues for periodic *reprocessing* of the Level 1 data set.

DMS-REQ-0313
DMS-REQ-0325

We plan to reprocess all image differencing-derived data (the Level 1 database), at the same time we perform the annual Level 2 data release productions. This will include all images taken since the start of survey operations, to the time when the data release production begins. The images will be reprocessed using a single version of the image differencing and measurement software, resulting in a consistent data set.

⁵⁴When Solar System objects are concerned, precovery has a slightly different meaning: predicting the positions of newly identified SSobjects on previously acquired visits, and associating with them the DIASources consistent with these predictions.

⁵⁵We will be maintaining a cache of 30 days of processed images to support this feature.

There will be three main advantages of Level 1 database produced during Data Release processing, compared to “living” Level 1 database: i) even the oldest data will be processed with the latest software, ii) astrometric and photometric calibration will be better, and iii) there will be no 12-month limit on the width of data window used to computed associated DIAObject measurements (proper motions, centroids, light curves, etc.).

diaCharacterization-
Cutoff

Older versions of the Level 1 database produced during Data Release processing will be archived following the same rules as for the Level 2 databases. The most recent DR, and the penultimate data release will be kept on disk and loaded into the database. Others will be archived to tape and made available as bulk downloads. See § 4.5 for more detail.

DMS-REQ-0077

DMS-REQ-0300

3.4 Level 1 Image Products

3.4.1 Visit Images

DMS-REQ-0069

Raw and Processed Visit Images will be made available for download no later than 24 hours from the end of visit acquisition.

L1PublicT

The images will remain accessible with low-latency (seconds from request to start of download) for at least 30 days, with slower access afterwards (minutes to hours).

I1CacheLifetime

3.4.2 Difference Images

DMS-REQ-0010

Complete difference images will be made available for download no later than 24 hours from the end of visit acquisition.

L1PublicT

The images will remain accessible with low-latency (seconds from request to start of download) for at least 30 days, with slower access afterwards (minutes to hours).

I1CacheLifetime

3.4.3 Image Differencing Templates

DMS-REQ-0280

Templates for difference image analysis will be created by coadding 6-months to a year long groups of visits. The coaddition process will take care to remove transient or fast moving objects (eg., asteroids) from the templates. Difference image analysis will use the appropriate template given the time of observation, airmass, and seeing.

3.5 Alerts to DIASources

3.5.1 Information Contained in Each Alert

For each detected DIASource, LSST will emit an “Event Alert” within 60 seconds of the end of visit (defined as the end of image readout from the LSST Camera). These alerts will be issued in V0Event format⁵⁶, and should be readable by V0Event-compliant clients.

OTT1

DMS-REQ-0002

Each alert (a V0Event packet) will at least include the following:

- *alertID*: An ID uniquely identifying this alert. It can also be used to execute a query against the Level 1 database as it existed when this alert was issued.
- *Level 1 database ID*
- Science Data:
 - The DIASource record that triggered the alert
 - The entire DIAObject (or SSObject) record
 - Previous 12 months of DIASource records
 - Matching Object IDs from the latest Data Release, if they exist, and 12 months of their DIASource records
- Cut-out of the difference image centered on the DIASource (10 bytes/pixel, FITS MEF)
- Cut-out of the template image centered on the DIASource (10 bytes/pixel, FITS MEF)

DMS-REQ-0274

diaCharacterization-Cutoff

The variable-size cutouts will be sized so as to encompass the entire footprint of the detected source, but be no smaller than 30×30 pixels. The provided images will comprise of a flux (32 bit float), variance (32 bit float), and mask (16 bit flags) planes, and include metadata necessary for further processing (e.g., WCS, zero-point, PSF, etc.).

The items above are meant to represent the *information* transmitted with each alert; the content of the alert packet itself will be formatted to confirm to V0Event (or other relevant) standard. Where the existing standard is inadequate for LSST needs, LSST will propose extensions and work with the community to reach a common solution.

⁵⁶Or some other format that is broadly accepted and used by the community at the start of LSST commissioning.

With each alert, we attempt to include as much information known to LSST about the DIA-Source as possible, to minimize the need for follow-up database queries. This speeds up classification and decision making at the user end, and relaxes the requirements on the database on the Project end.

3.5.2 Receiving and Filtering the Alerts

Alerts will be transmitted in VOEvent format, using standard IVOA protocols (e.g., VOEvent Transport Protocol; VTP⁵⁷). As a very high rate of alerts is expected, approaching ~10 million per night, we plan for public VOEvent Event Brokers⁵⁸ to be the primary end-points of LSST's event streams. End-users will use these brokers to classify and filter events for subsets fitting their science goals. End-users will *not* be able to subscribe to full, unfiltered, alert streams coming directly from LSST⁵⁹.

DMS-REQ-0002

To directly serve the end-users, LSST will provide a basic, limited capacity, alert filtering service. This service will run at the LSST U.S. Archive Center (at NCSA). It will let astronomers create simple filters that limit what alerts are ultimately forwarded to them⁶⁰. These *user defined filters* will be possible to specify using an SQL-like declarative language, or short snippets of (likely Python) code. For example, here's what a filter may look like:

DMS-REQ-0342

```
# Keep only never-before-seen events within two
# effective radii of a galaxy. This is for illustration
# only; the exact methods/members/APIs may change.
```

```
def filter(alert):
    if len(alert.sources) > 1:
        return False
    nn = alert.diaobject.nearest_neighbors[0]
    if not nn.flags.GALAXY:
        return False
```

⁵⁷VOEvent Transport Protocol is currently an IVOA Note, but we understand work is under way to finalize and bring it up to full IVOA Recommendation status.

⁵⁸These brokers are envisioned to be operated as a public service by third parties who will have signed MOUs with LSST.

⁵⁹This is due to finite network bandwidth available: for example, 100 end-users subscribing to a ~100 Mbps stream (the peak full stream data rate at end of the first year of operations) would require 10Gbps WAN connection from the archive center, just to serve the alerts.

⁶⁰More specifically, to their VTP clients. Typically, a user will use the Science User Interface (the web portal to LSST Archive Center) to set up the filters, and use their VTP client to receive the filtered VOEvent stream.

```
return nn.dist < 2. * nn.Re
```

We emphasize that this LSST-provided capability will be limited, and is *not* intended to satisfy the wide variety of use cases that a full-fledged public Event Broker could. For example, we do not plan to provide any *exclusive* classification to a unique category of object. Following the [SRD](#) specification, however, we will provide a limited number of pre-defined filters for a small number of object types of common interest. These will answer non-exclusive questions such as “is the light curve consistent with an RR Lyra?”, and will have potentially highly overlapping selections, designed to provide good completeness but perhaps only very modest purity. No information beyond what is contained in the `V0Event` packet will be available to the pre-defined or user-defined filters (e.g., no cross-matches to other catalogs). The complexity and run time of user defined filters will be limited by available resources. Execution latency will not be guaranteed. The number of `V0Events` transmitted to each user per visit will be limited as well (e.g., the equivalent of 20 full-size alert packets per visit per user, dynamically throttled depending on load). Finally, the total number of simultaneous subscribers is likely to be limited – in case of overwhelming interest, a TAC-like proposal process may be instituted.

DMS-REQ-0348

numBrokerAlerts
DMS-REQ-0343

numBrokerUsers

4 Level 2 Data Products

4.1 Overview

Level 2 data products result from direct image⁶¹ analysis. They're designed to enable *static* sky science (eg., studies of galaxy evolution, or weak lensing), and time-domain science that is not time sensitive (eg. statistical investigations of variability). They include image products (reduced single-epoch exposures, called *Processed Visit Images* or, sometimes, *calibrated exposures*, and coadds), and catalog products (tables of objects, sources, their measured properties, and related metadata).

DMS-REQ-0069
DMS-REQ-0279
DMS-REQ-0281
DMS-REQ-0275
DMS-REQ-0267
DMS-REQ-0276
DMS-REQ-0275

Similarly to Level 1 catalogs of DIAObjects and DIASources, Objects in the Level 2 catalog represent the astrophysical phenomena (stars, galaxies, quasars, etc.), while Sources represent their single-epoch observations. Sources are independently detected and measured in single epoch exposures and recorded in the Source table.

DMS-REQ-0267

The master list of Objects in Level 2 will be generated by associating and deblending the list of single-epoch DIASource detections and the lists of sources detected on coadds (Coadd-Sources). We plan to build coadds designed to maximize depth (*"deep coadds"*) and coadds designed to achieve a good combination of depth and seeing (*"best seeing coadds"*), unless algorithms will enable these two to be the same. We will also build a series of *short-period* (eg. yearly, or multi-year) coadds. The flux limit in deep coadds will be significantly fainter than in individual visits, and the best seeing coadds will help with deblending the detected sources. The short-period coadds are necessary to avoid missing faint objects showing long-term variability. These coadds will be built for all bands, as well as some combining multiple bands (*"multi-color coadds"*). **Not all of these will be preserved** after sources are detected and measured (see § 4.4.3 for details). We will provide a facility to regenerate their subsections as Level 3 tasks (§ 5).

DMS-REQ-0275
DMS-REQ-0279
DMS-REQ-0330
DMS-REQ-0337

DMS-REQ-0281

DMS-REQ-0311
DMS-REQ-0275

The deblender will be run simultaneously on the catalog of peaks⁶² detected in the coadds, the DIAObject catalog from the Level 1 database, and one or more external catalogs. It will use the knowledge of peak positions, bands, time, time variability (from Level 1 and the single-

⁶¹As opposed to *difference image*, for Level 1.

⁶²The source detection algorithm we plan to employ finds regions of connected pixels above the nominal S/N threshold in the *PSF-likelihood image* of the visit (or coadd). These regions are called *footprints*. Each footprint may have one or more *peaks*, and it is these peaks that the deblender will use to infer the number and positions of objects blended in each footprint.

epoch Source detections), inferred motion, Galactic longitude and latitude, and other available information to produce a master list of deblended Objects. Metadata on why and how a particular Object was deblended will be kept.

DMS-REQ-0276

The properties of Objects, including their exact positions, motions, parallaxes, and shapes, will be characterized by MultiFit-type algorithms⁶³.

Finally, to enable studies of variability, the fluxes of all Objects will be measured on individual visits (using both direct and difference images), with their shape parameters and deblending resolutions kept constant. This process is known as *forced photometry* (see § 4.2.4), and the flux measurements will be stored in the ForcedSource table.

DMS-REQ-0268

4.2 Level 2 Data Processing

Figures 3 and 4 present a high-level overview of the Level 2 data processing workflow⁶⁴. Logically⁶⁵, the processing begins with single-visit image reduction and source measurement, followed by global astrometric and photometric calibration, coadd creation, detection on coadds, association and deblending, object characterization, and forced photometry measurements.

The following is a high-level description of steps which will occur during regular Level 2 data processing (bullets 1 and 2 below map to pipeline 1, *Single Visit Processing*, in Figure 3, bullet 3 is pipeline 2, *Image Coaddition*, bullets 4-6 map to pipeline 3, *Coadded Image Analysis*, and bullet 7 is pipeline 4, *Multi-epoch Object Characterization*):

1. *Single Visit Processing*: Raw exposures are reduced to Processed Visit Images, and Sources are independently detected, deblended, and measured on all visits. Their measurements (instrumental fluxes and shapes) are stored in the Source table.
2. *Relative calibration*: The survey is internally calibrated, both photometrically and astrometrically. Relative zero-point and astrometric corrections are computed for every visit.

DMS-REQ-0267

DMS-REQ-0029

⁶³“MultiFit algorithms” are those that fit a PSF-convolved model (so-called “forward modeling”) to all multi-epoch observations of an object. This approach is in contrast to measurement techniques where multi-epoch images are coadded first, and the properties are measured from the coadded pixels.

⁶⁴Note that some LSST documents refer to *Data Release Processing*, which includes both Level 1 reprocessing (see § 3.3.5), and the Level 2 processing described here.

⁶⁵The actual implementation may parallelize these steps as much as possible.

Sufficient data is kept to reconstruct the normalized system response function $\phi_b(\lambda)$ (see Eq. 5, [SRD](#)) at every position in the focal plane at the time of each visit as required by § 3.3.4 of the [SRD](#).

DMS-REQ-0030

3. *Coadd creation*: Deep, seeing optimized, and short-period per-band coadds are created in *ugrizy* bands, as well as deeper, multi-color, coadds⁶⁶. Transient sources (including Solar System objects, explosive transients, etc), will be rejected from the coadds. See § 4.4.3 for details.

DMS-REQ-0279

DMS-REQ-0337

DMS-REQ-0281

4. *Coadd source detection*. Sources will be detected on all coadds generated in the previous step. The source detection algorithm will detect regions of connected pixels, known as *footprints*, above the nominal *S/N* threshold in the *PSF-likelihood image* of the visit. An appropriate algorithm will be run to also detect extended low surface brightness objects (eg., binned detection algorithm from SDSS). Each footprint may have one or more *peaks*, and the collection of these peaks (and their membership in the footprints) are the output of this stage.

DMS-REQ-0277

DMS-REQ-0349

5. *Association and deblending*. The next stage in the pipeline, which we will for simplicity just call *the deblender*, will synthesize a list of unique objects. In doing so it will consider the catalogs of CoaddSources, catalogs of DIASources, DIAObjects and SSObjects detected on difference images, and objects from external catalogs⁶⁷.

DMS-REQ-0275

DMS-REQ-0034

The deblender will make use of all information available at this stage, including the knowledge of peak positions, bands, time, time variability (from Level 1), Galactic longitude and latitude, etc. The output of this stage is a list of uncharacterized Objects⁶⁸.

6. *Coadd object characterization*. Object properties such as adaptive moments and aperture fluxes will be measured on the coadds. These will be stored in additional columns in the Object table. Models fit in multi-epoch object characterization will also be fit first to the coadds and stored.

DMS-REQ-0276

7. *Multi-epoch object characterization*. A set of predefined model fits and measurements will be performed on each of the Objects identified in the previous step, *taking all avail-*

⁶⁶We'll denote the "band" of the multi-color coadd as 'M'.

⁶⁷Note that Sources are not considered when generating the Object list (given the large number of visits in each band, the false positives close to the faint end would increase the complexity of association and deblending algorithms). It is possible for intermittent sources that are detected just above the faint detection limit of single visits to be undetected in coadds, and thus to not have a matching Object. To enable easy identification of such Sources, the nearest Object associated with each Source, if any, will be recorded.

⁶⁸Depending on the exact implementation of the deblender, this stage may also attach significant metadata (eg, deblended footprints and pixel-weight maps) to each deblended Object record.

able multi-epoch data into account. Model fits will be performed using *MultiFit*-type algorithms. Rather than coadding a set of images and measuring object characteristics on the coadd, MultiFit simultaneously fits PSF-convolved models to the objects multiple observations. This reduces systematic errors, improves the overall S/N , and allows for fitting of time-dependent quantities degenerate with shape on the coadds (for example, the proper motion). The models we plan to fit will *not* allow for flux variability (see the next item).

DMS-REQ-0275

8. *Forced Photometry*. Source fluxes will be measured on every visit, with the position, motion, shape, and the deblending parameters characterized in the previous step kept fixed. Measurements will be performed on both direct images and difference images. This process of *forced photometry*, will result in the characterization of the light-curve for each object in the survey. The fluxes will be stored in the ForcedSource table.

DMS-REQ-0268

4.2.1 Object Characterization Measures

Properties of detected objects will be measured as a part of the object characterization step described in the previous section and stored in the Object table. These measurements are designed to enable LSST “static sky” science. This section discusses at a high level which properties will be measured and how those measurements will be performed. For a detailed list of quantities being fit/measured, see the table in § 4.3.1.

All measurements discussed in this section deal with properties of *objects*, and will be performed on multi-epoch coadds, or by simultaneously fitting to all epochs. Measurements of sources in individual visits, independent of all others, are described in § 4.2.3.

To enable science cases depending on observations of non-variable objects in the LSST-observed sky, we plan to measure the following using the MultiFit approach:

- *Point source model fit*. The observed object is modeled as a point source with finite proper motion and parallax and constant flux (allowed to be different in each band). This model is a good description for non-variable stars and other unresolved sources. Its 11 parameters will be simultaneously constrained using information from all available observations in all bands⁶⁹.

DMS-REQ-0276

⁶⁹The fitting procedure will account for differential chromatic refraction.

- *Bulge-disk model fit.* The object is modeled as a sum of a de Vaucouleurs (Sersic $n = 4$) and an exponential (Sersic $n = 1$) component. This model is intended to be a reasonable description of galaxies⁷⁰. The object is assumed not to move⁷¹. The components share the same ellipticity and center. The model is independently fit to each band. There are a total of 8 free parameters, which will be simultaneously constrained using information from all available epochs for each band. Where there's insufficient data to constrain the likelihood (eg., small, poorly resolved, galaxies, or very few epochs), priors will be adopted to limit the range of its sampling.

DMS-REQ-0276

We will also explore fitting the galaxy model simultaneously to all bands, with some parameters constrained to be the same or related via a hierarchical model across bands. As this reduces the number of overall model parameters significantly, we could then consider freeing up other parameters. One likely scenario is that we would allow the bulge and disk ellipticities to differ; another possibility is allowing the Sersic indices of one or both components to vary. The ultimate determination of which model to use will be driven by empirical tests of the robustness and quality of the fits on both low- and high-redshift galaxies.

DMS-REQ-0333

In addition to the maximum likelihood values of fitted parameters and their covariance matrix, a number (currently planned to be ~ 200 , on average⁷²) of independent samples from the likelihood function will be provided. These will enable use-cases sensitive to departures from the Gaussian approximation, with shear measurement being the primary use case. A permissible descope, in case of insufficient storage, will be not to sample the posterior for u and y bands.

- *Standard colors.* Colors of the object in "standard seeing" (for example, the third quartile expected survey seeing in the i band, ~ 0.9 arcsec) will be measured. These colors are guaranteed to be seeing-insensitive, suitable for estimation of photometric redshifts⁷³.
- *Centroids.* Centroids will be computed independently for each band using an algorithm similar to that employed by SDSS. Information from all⁷⁴ epochs will be used to derive the estimate. These centroids will be used for adaptive moment, Petrosian, Kron, standard color, and aperture measurements.

DMS-REQ-0276

DMS-REQ-0276

⁷⁰We may reconsider this choice if a better suited parametrization is discovered while LSST is in Construction.

⁷¹I.e., have zero proper motion and trigonometric parallax.

⁷²This choice of the number of independent samples will be verified during Construction.

⁷³The problem of optimal determination of photometric redshift is the subject of intense research. The approach we're taking here is conservative, following contemporary practices. As new insights develop, we will revisit the issue.

⁷⁴Whenever we say *all*, it should be understood that this does not preclude reasonable data quality cuts to exclude data that would otherwise degrade the measurement.

- *Adaptive moments.* Adaptive moments will be computed using information from all epochs, independently for each band. The moments of the PSF realized at the position of the object will be provided as well. DMS-REQ-0276
- *Petrosian and Kron fluxes.* Petrosian and Kron radii and fluxes will be measured in standard seeing using self-similar elliptical apertures computed from adaptive moments. The apertures will be PSF-corrected and *homogenized*, convolved to a canonical circular PSF⁷⁵. The radii will be computed independently for each band. Fluxes will be computed in each band, by integrating the light within some multiple of *the radius measured in the canonical band*⁷⁶ (most likely the *i* band). Radii enclosing 50% and 90% of light will be provided. DMS-REQ-0276
- *Aperture surface brightness.* Aperture surface brightness will be computed in a variable number⁷⁷ of concentric, logarithmically spaced, PSF-homogenized, elliptical apertures, in standard seeing. DMS-REQ-0276
- *Variability characterization.* Parameters will be provided, designed to characterize periodic and aperiodic variability features [12], in each bandpass. We caution that the exact features in use when LSST begins operations are likely to be different compared to the baseline described here; this is to be expected given the rapid pace of research in time domain astronomy. However, their *number* is unlikely to grow beyond the present estimate. DMS-REQ-0276

4.2.2 Supporting Science Cases Requiring Full Posteriors

Science cases sensitive to systematics, departures of likelihood from Gaussianity, or requiring user-specified priors, demand knowledge of the shape of the likelihood function beyond a simple Gaussian approximation around the ML value. The estimate of bulge-disk model parameters and the estimate of photometric redshifts are two examples where knowledge of the full posterior is likely to be needed for LSST science cases. DMS-REQ-0046
DMS-REQ-0276

⁷⁵This is an attempt to derive a definition of elliptical apertures that does not depend on seeing. For example, for a large galaxy, the correction to standard seeing will introduce little change to measured ellipticity. Corrected apertures for small galaxies will tend to be circular (due to smearing by the PSF). In the intermediate regime, this method results in derived apertures that are relatively seeing-independent. Note that this is only the case for *apertures*; the measured flux will still be seeing dependent and it is up to the user to take this into account.

⁷⁶The shape of the aperture in all bands will be set by the profile of the galaxy in the canonical band alone. This procedure ensures that the color measured by comparing the flux in different bands is measured through a consistent aperture. See <http://www.sdss.org/dr7/algorithms/photometry.html> for details.

⁷⁷The number will depend on the size of the source.

We currently plan to provide this information in two ways: a) by providing independent samples from the likelihood function (in the case of the bulge-disk model), and b) by providing parametric estimates of the likelihood function (for the photometric redshifts). As will be shown in Table 4, the current allocation is ~ 200 samples (on average) for the bulge-disk model, and ~ 100 parameters for describing the photo-Z likelihood distributions, per object.

DMS-REQ-0333

The methods of storing likelihood functions (or samples thereof) will continue to be developed and optimized throughout Construction and Commissioning. The key limitation, on the amount of data needed to be stored, can be overcome by compression techniques. For example, simply noticing that not more than $\sim 0.5\%$ accuracy is needed for sample values allows one to increase the number of samples by a factor of 4. More advanced techniques, such as PCA analysis of the likelihoods across the entire catalog, may allow us to store even more, providing a better estimate of the shape of the likelihood function. In that sense, what is presented in Table 4 should be thought of as a **conservative estimate**, which we plan to improve upon as development continues in Construction.

4.2.3 Source Characterization

Sources will be detected on individual visits as well as the coadds. Sources detected on coadds will primarily serve as inputs to the construction of the master object list as described in § 4.2, and may support other LSST science cases as seen fit by the users (for example, searches for objects whose shapes vary over time).

DMS-REQ-0277

DMS-REQ-0267

The following Source properties are planned to be measured:

- *Static point source model fit.* The source is modeled as a static point source. There are a total of 3 free parameters (α , δ , flux). This model is a good description of stars and other unresolved sources.
- *Centroids.* Centroids are currently planned to be computed using an algorithm similar to that employed by SDSS. These centroids will be used for adaptive moment and aperture magnitude measurements.
- *Adaptive moments.* Adaptive moments will be computed. The moments of the PSF realized at the position of the object will be provided as well.
- *Aperture surface brightness.* Aperture surface brightness will be computed in a variable

DMS-REQ-0267

DMS-REQ-0267

DMS-REQ-0267

number⁷⁸ of concentric, logarithmically spaced, PSF-homogenized, elliptical apertures.

DMS-REQ-0267

Note that we do *not* plan to fit extended source Bulge+Disk models to individual Sources, nor measure per-visit Petrosian or Kron fluxes. These are object properties that are not expected to vary in time⁷⁹, and will be better characterized by MultiFit (in the Object table). For example, although a simple extendedness characterization is present in the Source table, star-galaxy separation (an estimate of the probability that a source is resolved, given the PSF) will be better characterized by MultiFit.

4.2.4 Forced Photometry

DMS-REQ-0268

DMS-REQ-0287

Forced Photometry is the measurement of flux in individual visits, given a fixed position, shape, and the deblending parameters of an object. It enables the study of time variability of an object's flux, irrespective of whether the flux in any given individual visit is above or below the single-visit detection threshold.

Forced photometry will be performed on all visits, for all Objects, using both direct images and difference images. The measured fluxes will be stored in the ForcedSources table. Due to space constraints, we only plan to measure the PSF flux.

4.2.5 Crowded Field Photometry

A fraction of LSST imaging will cover areas of high object (mostly stellar) density. These include the Galactic plane, the Large and Small Magellanic Clouds, and a number of globular clusters (among others).

LSST image processing and measurement software, although primarily designed to operate in non-crowded regions, is expected to perform well in areas of crowding. The current LSST applications development plan envisions making the deblender aware of Galactic longitude and latitude, and permitting it to use that information as a prior when deciding how to deblend objects. While not guaranteed to reach the accuracy or completeness of purpose-built crowded field photometry codes, we expect this approach will yield acceptable results even

⁷⁸The number will depend on the size of the source.

⁷⁹Objects that *do* change shape with time would, obviously, be of particular interest. Aperture fluxes provided in the Source table should suffice to detect these. Further per-visit shape characterization can be performed as a Level 3 task.

in areas of moderately high crowding.

Note that this discussion only pertains to processing of *direct images*. Crowding is not expected to significantly impact the quality of data products derived from *difference images* (i.e., Level 1).

4.3 The Level 2 Catalogs

This section presents the contents of key Level 2 catalog tables. As was the case for Level 1 (see § 3.3), here we present the *conceptual schemas* for the most important Level 2 tables (the Object, Source, and ForcedSource tables). The tables themselves are defined in LDM-153⁸⁰.

These convey *what* data will be recorded in each table, rather than the details of *how*. For example, columns whose type is an array (eg., radec) may be expanded to one table column per element of the array (eg., ra, dec1) once this schema is translated to SQL. Secondly, the tables to be presented are normalized (i.e., contain no redundant information). For example, since the band of observation can be found by joining a Source table to the table with exposure metadata, there's no column named band in the Source table. In the as-built database, the views presented to the users will be appropriately denormalized for ease of use.

DMS-REQ-0332

4.3.1 The Object Table

DMS-REQ-0275

Table 4: Level 2 Catalog Object Table

Name	Type	Unit	Description
objectId	uint64		Unique object identifier
parentObjectId	uint64		ID of the parent Object this object has been deblended from, if any.
radec	double[6][2]	arcsec	Position of the object (centroid), computed independently in each band. The centroid will be computed using an algorithm similar to that employed by SDSS.
radecErr	double[6][2]	arcsec	Uncertainty of radec.

Continued on next page

⁸⁰The SQL definition itself can be found in the cat package.

Table 4: Level 2 Catalog Object Table

Name	Type	Unit	Description
psRadecTai	double	time	Point source model: Time at which the object was at position psRadec.
psRadec	double[2]	degrees	Point source model: (α, δ) position of the object at time psRadecTai.
psPm	float[2]	mas/yr	Point source model: Proper motion vector.
psParallax	float	mas	Point source model: Parallax.
psFlux	float[ugrizy]	nmgy	Point source model fluxes ⁸¹ .
psCov	float[66]	various	Point-source model covariance matrix ⁸² .
psLnL	float		Natural $\log$ likelihood of the observed data given the point source model.
psChi2	float		χ^2 statistic of the model fit.
psNdata	int		The number of data points (pixels) used to fit the model.
bdRadec	double[2][ugrizy]		B+D model ⁸³ : (α, δ) position of the object, in each band.
bdEllip	float[2][ugrizy]		B+D model: Ellipticity (e_1, e_2) of the object.
bdFluxB	float[ugrizy]	nmgy	B+D model: Integrated flux of the de Vaucouleurs component.
bdFluxD	float[ugrizy]	nmgy	B+D model: Integrated flux of the exponential component.
bdReB	float[ugrizy]	arcsec	B+D model: Effective radius of the de Vaucouleurs profile component.
bdReD	float[ugrizy]	arcsec	B+D model: Effective radius of the exponential profile component.

Continued on next page

⁸¹Point source model assumes that fluxes are constant in each band. If the object is variable, psFlux will effectively be some estimate of the average flux.

⁸²Not all elements of the covariance matrix need to be stored with same precision. While the variances will be stored as 32 bit floats ($\sim$ seven significant digits), the covariances may be stored to $\sim$ three significant digits ($\sim 1\%$).

⁸³Though we refer to this model as “Bulge plus Disk”, we caution the reader that the decomposition, while physically motivated, should not be taken too literally.

Table 4: Level 2 Catalog Object Table

Name	Type	Unit	Description
bdCov	float[36][ugrizy]		B+D model covariance matrix ⁸⁴ .
bdLnL	float[ugrizy]		Natural <i>log</i> likelihood of the observed data given the bulge+disk model.
bdChi2	float[ugrizy]		χ^2 statistic of the model fit.
bdNdata	int[ugrizy]		The number of data points (pixels) used to fit the model.
bdSamples	float16[9][200][ugrizy]		Independent samples of bulge+disk likelihood surface. All sampled quantities will be stored with at least ~ 3 significant digits of precision. The number of samples will vary from object to object, depending on how well the object's likelihood function is approximated by a Gaussian.
stdColor	float[5]	mag	Color of the object measured in "standard seeing". While the exact algorithm is yet to be determined, this color is guaranteed to be seeing-independent and suitable for photo-Z determinations.
stdColorErr	float[5]	mag	Uncertainty of stdColor.
lxx	float	asec ²	Adaptive second moment of the source intensity. See Bernstein & Jarvis [2] for detailed discussion of all adaptive-moment related quantities ⁸⁵ .
lyy	float	asec ²	Adaptive second moment of the source intensity.
lxy	float	asec ²	Adaptive second moment of the source intensity.

Continued on next page

⁸⁴See psCov for notes on precision of variances/covariances.

⁸⁵Or <http://ls.st/5f4> for a brief summary.

Table 4: Level 2 Catalog Object Table

Name	Type	Unit	Description
lcov	float[6]	asec ⁴	Ixx, Iyy, Ixy covariance matrix.
lxxPSF	float	asec ²	Adaptive second moment for the PSF.
lyyPSF	float	asec ²	Adaptive second moment for the PSF.
lxyPSF	float	asec ²	Adaptive second moment for the PSF.
m4	float[ugrizy]		Fourth order adaptive moment.
petroRad	float[ugrizy]	arcsec	Petrosian radius, computed using elliptical apertures defined by the adaptive moments.
petroRadErr	float[ugrizy]	arcsec	Uncertainty of petroRad
petroBand	int8		The band of the canonical petroRad
petroFlux	float[ugrizy]	nmgy	Petrosian flux within a defined multiple of the canonical petroRad
petroFluxErr	float[ugrizy]	nmgy	Uncertainty in petroFlux
petroRad50	float[ugrizy]	arcsec	Radius containing 50% of Petrosian flux.
petroRad50Err	float[ugrizy]	arcsec	Uncertainty of petroRad50.
petroRad90	float[ugrizy]	arcsec	Radius containing 90% of Petrosian flux.
petroRad90Err	float[ugrizy]	arcsec	Uncertainty of petroRad90.
kronRad	float[ugrizy]	arcsec	Kron radius (computed using elliptical apertures defined by the adaptive moments)
kronRadErr	float[ugrizy]	arcsec	Uncertainty of kronRad
kronBand	int8		The band of the canonical kronRad
kronFlux	float[ugrizy]	nmgy	Kron flux within a defined multiple of the canonical kronRad
kronFluxErr	float[ugrizy]	nmgy	Uncertainty in kronFlux
kronRad50	float[ugrizy]	arcsec	Radius containing 50% of Kron flux.
kronRad50Err	float[ugrizy]	arcsec	Uncertainty of kronRad50.
kronRad90	float[ugrizy]	arcsec	Radius containing 90% of Kron flux.
kronRad90Err	float[ugrizy]	arcsec	Uncertainty of kronRad90.

Continued on next page

Table 4: Level 2 Catalog Object Table

Name	Type	Unit	Description
apNann	int8		Number of elliptical annuli (see below).
apMeanSb	float[6][apNann]	nmgy/as ²	Mean surface brightness within an annulus ⁸⁶ .
apMeanSbSigma	float[6][apNann]	nmgy/as ²	Standard deviation of apMeanSb.
extendedness	float		A measure of extendedness, computed using a combination of available moments, or from a likelihood ratio of point/B+D source models (exact algorithm TBD). <i>extendedness</i> = 1 implies a high degree of confidence that the source is extended. <i>extendedness</i> = 0 implies a high degree of confidence that the source is point-like.
lcPeriodic	float[6 × 32]		Periodic features extracted from difference image-based light-curves using generalized Lomb-Scargle periodogram [Table 4, 12].
lcNonPeriodic	float[6 × 20]		Non-periodic features extracted from difference image-based light-curves [Table 5, 12].
photoZ	float[2 × 100]		Photometric redshift likelihood samples – pairs of (<i>z</i> , <i>logL</i>) – computed using a to-be-determined published and widely accepted algorithm at the time of LSST Commissioning.
flags	bit[128]	bit	Various useful flags.

⁸⁶A database function will be provided to compute the area of each annulus, to enable the computation of aperture flux.

4.3.2 Source Table

Source measurements are performed independently on individual visits. They're designed to enable relative astrometric and photometric calibration, variability studies of high signal-to-noise objects, and studies of high SNR objects that vary in position and/or shape (eg., comets).

DMS-REQ-0267

Table 5: Level 2 Catalog Source Table

Name	Type	Unit	Description
sourceId	uint64		Unique source identifier ⁸⁷
ccdVisitId	uint64		ID of CCD and visit where this source was measured
objectId	uint64		ID of the object this source was associated with, if any.
ssObjectId	uint64		ID of the SSObject this source has been linked to, if any.
parentSourceId	uint64		ID of the parent Source this source has been deblended from, if any.
xy	float[2]	pixels	Position of the object (centroid), computed using an algorithm similar to that used by SDSS.
xyCov	float[3]		Covariance matrix for xy.
radec	double[2]	arcsec	Calibrated (α, δ) of the source, transformed from xy.
radecCov	float[3]	arcsec	Covariance matrix for radec.
apFlux	float	nmgy	Calibrated aperture flux.
apFluxErr	float	nmgy	Estimated uncertainty of apFlux.
sky	float	nmgy/asec ²	Estimated background (sky) surface brightness at the position (centroid) of the source.
skyErr	float	nmgy/asec ²	Estimated uncertainty of sky.
psRadec	double[2]	degrees	Point source model: (α, δ) position of the object.
psFlux	float	nmgy	Calibrated point source model flux.

Continued on next page

⁸⁷It would be optimal if the source ID is globally unique across all releases. Whether that's realized will depend on technological and space constraints.

Table 5: Level 2 Catalog Source Table

Name	Type	Unit	Description
psCov	float[6]	various	Point-source model covariance matrix ⁸⁸ .
psLnL	float		Natural $\log$ likelihood of the observed data given the point source model.
psChi2	float		χ^2 statistic of the model fit.
psNdata	int		The number of data points (pixels) used to fit the model.
lxx	float	asec ²	Adaptive second moment of the source intensity. See Bernstein & Jarvis [2] for detailed discussion of all adaptive-moment related quantities ⁸⁹ .
lyy	float	asec ²	Adaptive second moment of the source intensity.
lxy	float	asec ²	Adaptive second moment of the source intensity.
lcov	float[6]	asec ⁴	Ixx, Iyy, Ixy covariance matrix.
lxxPSF	float	asec ²	Adaptive second moment for the PSF.
lyyPSF	float	asec ²	Adaptive second moment for the PSF.
lxyPSF	float	asec ²	Adaptive second moment for the PSF.
apNann	int8		Number of elliptical annuli (see below).
apMeanSb	float[apNann]	nmgy	Mean surface brightness within an annulus.
apMeanSbSigma	float[apNann]	nmgy	Standard deviation of apMeanSb.

Continued on next page

⁸⁸Not all elements of the covariance matrix will be stored with same precision. While the variances will be stored as 32 bit floats (~ seven significant digits), the covariances may be stored to ~ three significant digits (~1%).

⁸⁹Or <http://ls.st/5f4> for a brief summary.

Table 5: Level 2 Catalog Source Table

Name	Type	Unit	Description
extendedness	float		A measure of extendedness, computed using a combination of available moments (exact algorithm TBD). <i>extendedness</i> = 1 implies a high degree of confidence that the source is extended. <i>extendedness</i> = 0 implies a high degree of confidence that the source is point-like.
flags	bit[64]	bit	Various useful flags.

4.3.3 ForcedSource Table

DMS-REQ-0268

Table 6: Level 2 Catalog ForcedSource Table

Name	Type	Unit	Description
objectId	uint64		Unique object identifier
ccdVisitId	uint64		ID of CCD and visit where this source was measured
psFlux	float	nmgy	Point source model flux on direct image, if performed.
psFluxErr	float	nmgy	Point source model flux error, stored to 1% precision.
psDiffFlux	float	nmgy	Point source model flux on difference image, if performed.
psDiffFluxErr	float	nmgy	Point source model flux error, stored to 1% precision.
flags	bit[8]	bit	Various useful flags.

4.4 Level 2 Image Products

4.4.1 Visit Images

DMS-REQ-0065

Raw exposures, including individual snaps, and Processed Visit Images will be made available for download as FITS files. They will be downloadable both through a human-friendly Science User Interface, as well as using machine-friendly APIs.

Required calibration data, processing metadata, and all necessary image processing software will be provided to enable the user to generate bitwise identical processed images from raw images⁹⁰.

DMS-REQ-0130

DMS-REQ-0298

DMS-REQ-0308

4.4.2 Calibration Data

All calibration frames (darks, flats, biases, fringe, etc.) will be preserved and made available for download as FITS files.

DMS-REQ-0130

DMS-REQ-0065

DMS-REQ-0298

All auxiliary telescope data, both raw (images with spectra) and processed (calibrated spectra, derived atmosphere models), will be preserved and made available for download.

4.4.3 Coadded Images

In course of Level 2 processing, multiple classes and numerous of coadds will be created⁹¹:

- A set of *deep coadds*. One deep coadd will be created for each of the *ugrizy* bands, plus a seventh, deeper, multi-color coadd. These coadds will be optimized for a reasonable combination of depth (i.e., employ no PSF matching) and resolution (i.e., visits with significantly degraded seeing may be omitted). Transient sources (including Solar System objects, explosive transients, etc), will be removed. Care will be taken to preserve the astrophysical backgrounds⁹².

DMS-REQ-0279

DMS-REQ-0281

The six per-band coadds will be kept indefinitely and made available to the users. *Their primary purpose is to enable the end-users to apply alternative object characterization algorithms, perform studies of diffuse structures, and for visualization.*

DMS-REQ-0334

⁹⁰Assuming identically performing software and hardware configuration.

⁹¹See LDM-151 for more details.

⁹²For example, using “background matching” techniques [9]

- A set of *best seeing* coadds. One deep coadd will be created for each of the *ugrizy* bands, using only the best seeing data (for example, using only the first quartile of the realized seeing distribution). These will be built to assist the deblending process. These coadds will be kept indefinitely and made available to the users. DMS-REQ-0330
DMS-REQ-0334
- A set of *short-period* coadds. These will comprise of multiple (*ugrizyM*) sets of yearly and multi-year coadds. Each of these sets will be created using only a subset of the data, and otherwise share the characteristics of the deep coadds described above. These are designed to enable detection of long-term variable or moving⁹³ objects that would be “washed out” (or rejected) in full-depth coadds. **We do not plan to keep and make these coadds available.** We will retain and provide sufficient metadata for users to re-create them using Level 3 or other resources. DMS-REQ-0337
DMS-REQ-0298
DMS-REQ-0311
- One (*ugrizyM*) set of PSF-matched coadds. These will be used to measure colors and shapes of objects at “standard” seeing. **We do not plan to keep and make these coadds available.** We will retain and provide sufficient metadata for users to re-create them using Level 3 or other resources. DMS-REQ-0335

The exact details of which coadds to build, and which ones to keep, can change during Construction without affecting the processing system design as the most expensive operations (raw image input and warping) are constant in the number of coadds produced. The data management system design *is* sensitive to the total *number and size* of coadds to be *kept* – these are the relevant constraining variables.

We reiterate that **not all coadds will be kept and served to the public**⁹⁴, though sufficient metadata will be provided to users to recreate them on their own. Some coadds may be entirely “virtual”: for example, the PSF-matched coadds could be implemented as ad-hoc convolutions of postage stamps when the colors are measured. DMS-REQ-0311

We will retain smaller sections of all generated coadds, to support quality assessment and targeted science. Retained sections may be positioned to cover areas of the sky of special interest such as overlaps with other surveys, nearby galaxies, large clusters, etc. DMS-REQ-0338

⁹³For example, nearby high proper motion stars.

⁹⁴The coadds are a major cost driver for storage. LSST Data Management system is currently sized to keep and serve seven coadds, *ugrizyM*, over the full footprint of the survey.

4.5 Data Release Availability and Retention Policies

Over 10 years of operations, LSST will produce eleven data releases: two for the first year of survey operations, and one every subsequent year. Each data release will include reprocessing of all data from the start of the survey, up to the cutoff date for that release.

The contents of data releases are expected to range from a few PB (DR1) to ~70 PB for DR11 (this includes the raw images, retained coadds, and catalogs). Given that scale, it is not feasible to keep all data releases loaded and accessible at all times.

Instead, ***only the contents of the most recent data release, and the penultimate data release will be kept on fast storage and with catalogs loaded into the database.*** Statistics collected by prior surveys (eg., SDSS) show that users nearly always prefer accessing the most recent data release, but sometimes may use the penultimate one (this is especially true just after the publication of a new data release). Older releases are used rarely.

DMS-REQ-0313

To assist with data quality monitoring and assessment ***small, overlapping, samples of data from older releases will be kept loaded in the database.*** The sample size is expected to be on order of ~1-5% of the data release data, with larger samples kept early on in the survey. The goal is to allow one to test how the reported characterization of the same data varies from release to release.

DMS-REQ-0339

Older releases will be archived to mass storage (tape). The users ***will not be able to perform database queries against archived releases.*** They will be made available as bulk downloads in some common format (for example, FITS binary tables). Database software and data loading scripts will be provided for users who wish to set up a running copy of an older (or current) data release database on their systems.

DMS-REQ-0077

DMS-REQ-0300

All raw data used to generate any public data product (raw exposures, calibration frames, telemetry, configuration metadata, etc.) will be kept and made available for download.

DMS-REQ-0309

DMS-REQ-0346

5 Level 3 Data Products and Capabilities

Level 3 capabilities are envisioned to enable science cases that would greatly benefit from co-location of user processing and/or data within the LSST Archive Center. The high-level requirement for Level 3 is established in § 3.5 of the LSST [SRD](#).

OSS-REQ-0140

Level 3 capabilities include three separate deliverables:

1. Level 3 Data Products and associated storage resources
2. Level 3 processing resources, and
3. Level 3 programming environment and framework

Many scientists' work may involve using two or all three of them in concert, but they can each be used independently. We describe each one of them in the subsections to follow.

5.1 Level 3 Data Products and Associated Storage Resources

DMS-REQ-0290

These are data products that are generated by users *on any computing resources anywhere* that are then brought to an LSST Data Access Center (DAC) and stored there. The hardware for these capabilities includes the physical storage and database server resources at the DAC to support them.

For catalog data products, there is an expectation that they can be “federated” with the Level 1 (L1) and Level 2 (L2) catalogs to enable analyses combining them. Essentially this means that either the user-supplied tables include keys from the L1/L2 catalogs that can be used for key-equality-based joins with them (example: a table of custom photometric redshifts for galaxies, with a column of object IDs that can be joined to the L2 Object catalog), or that there are columns that can be used for spatial (or temporal, or analogous) joins against L1/L2 tables. The latter implies that such L3 table's columns must be in the same coordinate system and units as the corresponding L1/L2 columns.

DMS-REQ-0123

DMS-REQ-0124

There is no requirement that Level 3 data products (L3DPs) are derived from L1 or L2 other than that they be joinable with them. For instance, a user might have a catalog of radio sources that they might want to bring into federation with the LSST catalogs. That can be

thought of as a Level 3 Data Product as long as they have “LSST-ized” it by ensuring compatibility of coordinate, time, measurement systems, etc. Nevertheless, we do expect the majority of L3DPs to be derived from processed LSST data.

There could also be L3 image data products; for example, user-generated coadds with special selection criteria or stacking algorithms (eg. the so-called shift & stack algorithm for detecting moving objects).

DMS-REQ-0127

Any L3DP may have access controls associated with it, restricting read access to just the owner, to a list of people, to a named group of people, or allowing open access.

DMS-REQ-0340

The storage resources for L3DPs come out of the [SRD](#) requirement for 10% of LSST data management capabilities to be devoted to user processing. In general, they are likely to be controlled by some form of a “space allocation committee”. Users will probably have some small baseline automatic allocation, beyond which a SAC proposal is needed. The SAC may take into account scientific merit, length of time for which the storage is requested, and openness of the data to others, in setting its priorities.

DMS-REQ-0119

It is to be decided whether users will be required to provide the code and/or documentation behind their L3DPs, or whether the SAC may include the availability of this supporting information in its prioritization. Obviously if a user intends to make a L3DP public or publish it to a group it will be more important that supporting information be available.

Level 3 data products that are found to be generally useful can be migrated to Level 2. This is a fairly complex process that ultimately involves the project taking responsibility for supporting and running LSST-style code that implements the algorithm necessary to produce the data product (it’s not just relabeling an existing L3DP as L2). The project will provide necessary support for such migrations.

5.2 Level 3 Processing Resources

DMS-REQ-0119

These are project-owned computing resources located at the DACs. They are available for allocation to all users with LSST data rights. They may be used for any computation that involves the LSST data and advances LSST-related science. The distinctive feature of these computing resources is that they are located with excellent I/O connections to the image and catalog datasets at Level 1 and Level 2. There may be other co-located but *not* project-owned,

resources available at the LSST DACs⁹⁵; their use is beyond the scope of this document, except to note that reasonable provisions will be made to ensure it is possible to use them to process large quantities of LSST data.

Level 3 processing resources will, at least, include systems that can carry out traditional batch-style processing, probably similarly configured to those LSST will be using for the bulk of data release production processing. It is to be determined whether any other flavors of hardware would be provided, such as large-memory machines; this is likely to be driven by the project need (or lack thereof) for such resources.

There will be a time allocation committee (TAC) for these resources. Every LSST-data-rights user may get a small default allocation (enough to run test jobs). Substantial allocations will require a scientific justification. Priorities will be based on the science case and, perhaps, also on whether the results of the processing will be released to a larger audience. Requests must specify what special flavors of computing will be needed (e.g., GPUs, large memory, etc.).

A fairly standard job control environment (like Condor), will be available, and users will be permitted to work with it at a low, generic level. They will not be required to use the higher levels of the LSST process control middleware (but they may; see § 5.3).

These processing resources can be available for use in any clearly LSST-related scientific work. It is not strictly required that they be used to process LSST data, in this context. For instance, it could be acceptable to run special types of cosmological simulations that are in direct support of an LSST analysis, *if the closeness to the data makes the LSST facility uniquely suitable for such work*. The TAC will take into account in its decisions whether proposed work makes good use of the enhanced I/O bandwidth available to LSST data on these systems.

5.3 Level 3 Programming Environment and Framework

DMS-REQ-0125

DMS-REQ-0128

As a part of the Level 3 Programming Environment and Framework, the LSST will make available the LSST software stack to users, to aid in the analyses of LSST data. This includes all code implementing the core processing algorithms (image characterization and manipulation, building of coadds, image differencing, object detection, object characterization, moving object detection, etc.), the middleware necessary to run those codes at large scale, as well as

⁹⁵For example, the U.S. DAC will be located at the National Petascale Facility building at NCSA, adjacent to the Blue Waters supercomputer.

the LSST database management system.

DMS-REQ-0033
DMS-REQ-0032
DMS-REQ-0308

These analyses could be done on LSST-owned systems (i.e., on the Level 3 processing resources) but also on a variety of supported external systems. We will aim to support common personal Unix flavors (for example, common distributions of Linux and Mac OS X) as well as commonly used cluster and HPC environments. The vision is to enable relatively straightforward use of major national systems such as XSEDE or Open Science Grid, as well as some common commercial cloud environments. The decision of which environments to support will be under configuration control and we will seek advice from the user community. We cannot commit to too many flavors. In-kind contributions of customizations for other environments will be welcome and may provide a role for national labs.

The Level 3 environment is intended, when put to fullest use, to allow users to run their own productions-like runs on bulk image and/or catalog data, with mechanisms for creating and tracking large groups of jobs in a batch system.

The Level 3 environment, in asymptopia, has a great deal in common with the environment that the Project will use to build the Level 2 data releases. It is distinct, however, as supporting it as a tool meant for the end-users imposes additional requirements:

- In order to be successful as a *user* computing environment, it needs to be easy to use. Experience with prior projects⁹⁶ has shown that if the production computing environment is not envisioned from the start as being shared with users, it will likely evolve into an experts-only tool that is too complicated, or too work-hardened, to serve users well.
- While it is desirable for the production computing to be portable to Grid, cloud, etc. resources, this option might not be exercised in practice and could atrophy. For the user community, it's a far more central capability. Early community engagement is therefore key to developing and maintaining these capabilities.
- Not all the capabilities of the LSST production environment need necessarily be exported to the users. LSST-specific capabilities associated with system administration, for instance, are not of interest to end-users.

⁹⁶For example, BaBar.

5.4 Migration of Level 3 data products to Level 2

- For the migration to be considered, the creator of the L3DP will need to agree to make their data product public to the entire LSST data-rights community, along with supporting documentation and code. The code at first need not be in the LSST framework or even in an LSST-supported language.
- If the original proponent wrote her/his code in the C++/Python LSST stack environment (the “Level 3 environment”), it will be easier to migrate it to Level 2 (though, obviously, using the same languages/frameworks does not guarantee that the code is of production quality).
- If the original code was written in another language or another data processing framework, the project may consider rewriting it to required LSST standards.
- Taking on a new Level 2 DP means that the project is committing to code maintenance, data quality review, space allocation, and continuing production of the new L2DP through DR11.

6 Data Products for Special Programs

LSR-REQ-0075

LSST Survey Specifications (LSST [SRD](#), § 3.4) specify that 90% of LSST observing time will be spend executing the so-called “universal cadence”. These observations will result in Level 1 and 2 data products described earlier in this document.

The remaining 10% of observing time will be devoted to special programs, obtaining improved coverage of interesting regions of observational parameter space. Examples include very deep ($r \sim 26$, per exposure) observations, observations with very short revisit times (~ 1 minute), and observations of “special” regions such as the Ecliptic, Galactic plane, and the Large and Small Magellanic Clouds. A third type of survey, micro-surveys, that would use about 1% of the time, may also be considered.

The details of these special programs or micro surveys are not yet defined⁹⁷. Consequently, the specifics of their data products are left undefined at this time. Instead, we just specify the *constraints* on these data products, given the adopted Level 1/2/3 architecture. It is understood that no special program will be selected that does not fit these constraints⁹⁸. This allows us to size and construct the data management system, without knowing the exact definition of these programs this far in advance.

Processing for special programs will make use of the same software stack and computing capabilities as the processing for universal cadence. The programs are expected to use no more than $\sim 10\%$ of computational and storage capacity of the LSST data processing cluster. When special products include time domain event alerts, their processing shall generally be subject to the same latency requirements as Level 1 data products.

DMS-REQ-0320

DMS-REQ-0321
DMS-REQ-0344
OTT1

For simplicity of use and consistency, the data products for special programs will be stored in databases separate from the “main” (Level 1 and 2) databases. The system will, however, allow for simple federation with Level 1/2/3 data products (i.e., cross-queries and joins).

DMS-REQ-0322

As a concrete example, a data product complement for a “deep drilling” field designed for supernova discovery and characterization may consist of: i) alerts to events discovered by differencing the science images against a special deep drilling template, ii) a Level 1-like database iii) one or more “nightly co-adds” (co-adds built using the data from the entire night), pro-

⁹⁷The initial complement is expected to be defined and selected no later than Science Verification.

⁹⁸Or will come with additional, external, funding, capabilities, and/or expertise.

duced and made available within ~24 hours, and iv) special deep templates, built using the best recently acquired seeing data, produced on a fortnightly cadence.

L1PublicT

Note that the data rights and access rules apply just as they would for for Level 1/2/3. For example, while generated event alerts (if any) will be accessible world-wide, the image and catalog products will be restricted to users with LSST data rights.

References

- [1] **[LDM-153]**, Becla, J., 2013, *LSST Database Baseline Schema*, LDM-153, URL <https://ls.st/LDM-153>
- [2] Bernstein, G.M., Jarvis, M., 2002, *AJ*, 123, 583 (arXiv:astro-ph/0107431), doi:10.1086/338085, ADS Link
- [3] **[LSE-29]**, Claver, C.F., The LSST Systems Engineering Integrated Project Team, 2016, *LSST System Requirements*, LSE-29, URL <https://ls.st/LSE-29>
- [4] **[LSE-30]**, Claver, C.F., The LSST Systems Engineering Integrated Project Team, 2016, *LSST System Requirements*, LSE-30, URL <https://ls.st/LSE-30>
- [5] **[LSE-61]**, Dubois-Felsmann, G., Jenness, T., 2017, *LSST Data Management Subsystem Requirements*, LSE-61, URL <https://ls.st/LSE-61>
- [6] **[LPM-17]**, Ivezić, Ž., The LSST Science Collaboration, 2011, *LSST Science Requirements Document*, LPM-17, URL <https://ls.st/LPM-17>
- [7] Ivezić, Z., et al., 2008, ArXiv e-prints (arXiv:0805.2366), ADS Link
- [8] **[LDM-133]**, Jurić, M., Lim, K.T., Kantor, J., 2013, *Data Management UML Domain Model*, LDM-133, URL <https://ls.st/LDM-133>
- [9] **[DMTN-035]**, Jurić, M., Becker, A., Shaw, R., Krughoff, K.S., Kantor, J., 2016, *Winter 2013 LSST DM Data Challenge Release Notes*, DMTN-035, URL <https://dmtn-035.lsst.io>, LSST Data Management Technical Note
- [10] Muinonen, K., Belskaya, I.N., Cellino, A., et al., 2010, *Icarus*, 209, 542, doi:10.1016/j.icarus.2010.04.003, ADS Link
- [11] **[LDM-156]**, Myers, J., Jones, L., Axelrod, T., 2013, *Moving Object Pipeline System Design*, LDM-156, URL <https://ls.st/LDM-156>

- [12] Richards, J.W., Starr, D.L., Butler, N.R., et al., 2011, *ApJ*, 733, 10 (arXiv:1101.1959), doi:10.1088/0004-637X/733/1/10, ADS Link
- [13] **[LDM-151]**, Swinbank, J.D., et al., 2017, *Data Management Science Pipelines Design*, LDM-151, URL <https://ls.st/LDM-151>

A Appendix: Conceptual Pipeline Design

A high-level conceptual overview of the LSST image processing science pipelines is illustrated in Figure 2. The pipeline definitions presented here are driven by their inputs, outputs and processing steps; they do *not* describe exact boundaries in the actual implementation code, execution, or development responsibilities within the Project. Processing from pipelines marked with 1, 2, and 5-8 is executed every day when new data are taken to produce Level 1 Data Products. Annual Data Release processing includes pipelines 1-6 and 8 (everything except Alert production). These main conceptual steps in LSST image processing include the following pipelines (enumeration in this list corresponds to enumeration in Figure 2 but note that these steps can be interleaved in the actual processing flow):

1. *Single Visit Processing* pipeline (Figure 3) produces calibrated and characterized single-visit images from raw snaps. The main processing steps include instrumental signature removal, background estimation, source detection, deblending and measurements, point spread function estimation, and astrometric and photometric calibration.
2. *Image Coaddition* pipeline (Figure 4) produces coadded images of different flavors (optimized for depth, seeing, etc.) from an ensemble of single-visit images.
3. *Coadded Image Analysis* pipeline (Figure 4) defines the Object list and performs initial measurements on coadded images.
4. *Multi-epoch Object Characterization* pipeline (Figure 4) fits a library of image models to a set of Footprints of an Object family, measures additional quantities (e.g., the profile surface brightness in a series of annuli) not captured by those models, and performs Forced Photometry. All these measurements are performed on single-visit images (direct and difference images) and for all Objects.
5. *Image Differencing* pipeline (Figure 5) produces difference images from a single-visit and coadded (template) images.
6. *Difference Image Analysis* pipeline (Figure 5) updates DIAObject and SSObject lists with new DIASources detected on processed difference image, fits a library of image models to Footprints of these DIASources, and for all DIAObjects overlapping the difference image it performs Forced Photometry and recomputes summary quantities. During nightly Level 1 processing, this pipeline also performs Forced Photometry for all new DIAObjects on difference images from the last 30 days.

```
precoveryWindow
```




FIGURE 2: Illustration of the conceptual design of LSST science pipelines for imaging processing.

7. *Alert Generation and Distribution* pipeline (Figure 5) uses updated DIAObjects and DIA-Sources to generate and distribute Alerts (which also include postage stamp images of the DIASource in difference image and coadded template image).
8. *Moving Object Processing* pipeline (MOPS, Figure 6) combines all un-associated DIASources into plausible SSObjects and estimates their orbital parameters. The three main pipeline stages include associating new DIASources with known SSObjects, discovering new SS-Objects, and orbit refinement and management.

Further details about the pipeline design and implementation are available from the LSST document LDM-151.




FIGURE 3: Illustration of the conceptual algorithm design for Single Visit Processing pipeline.






FIGURE 4: Illustration of the conceptual algorithm design for Image Coaddition, Coadded Image Analysis, and Multi-epoch Object Characterization pipelines.

5 Image Differencing



6 Difference Image Analysis



7 Alert Generation and Distribution




FIGURE 5: Illustration of the conceptual algorithm design for Image Differencing, Difference Image Analysis, and Alert Generation and Distribution pipelines.

8 Moving Object Processing System




FIGURE 6: Illustration of the conceptual algorithm design for the Moving Object Processing Software pipeline.