

Rubin Observatory

Vera C. Rubin Observatory
Data Management

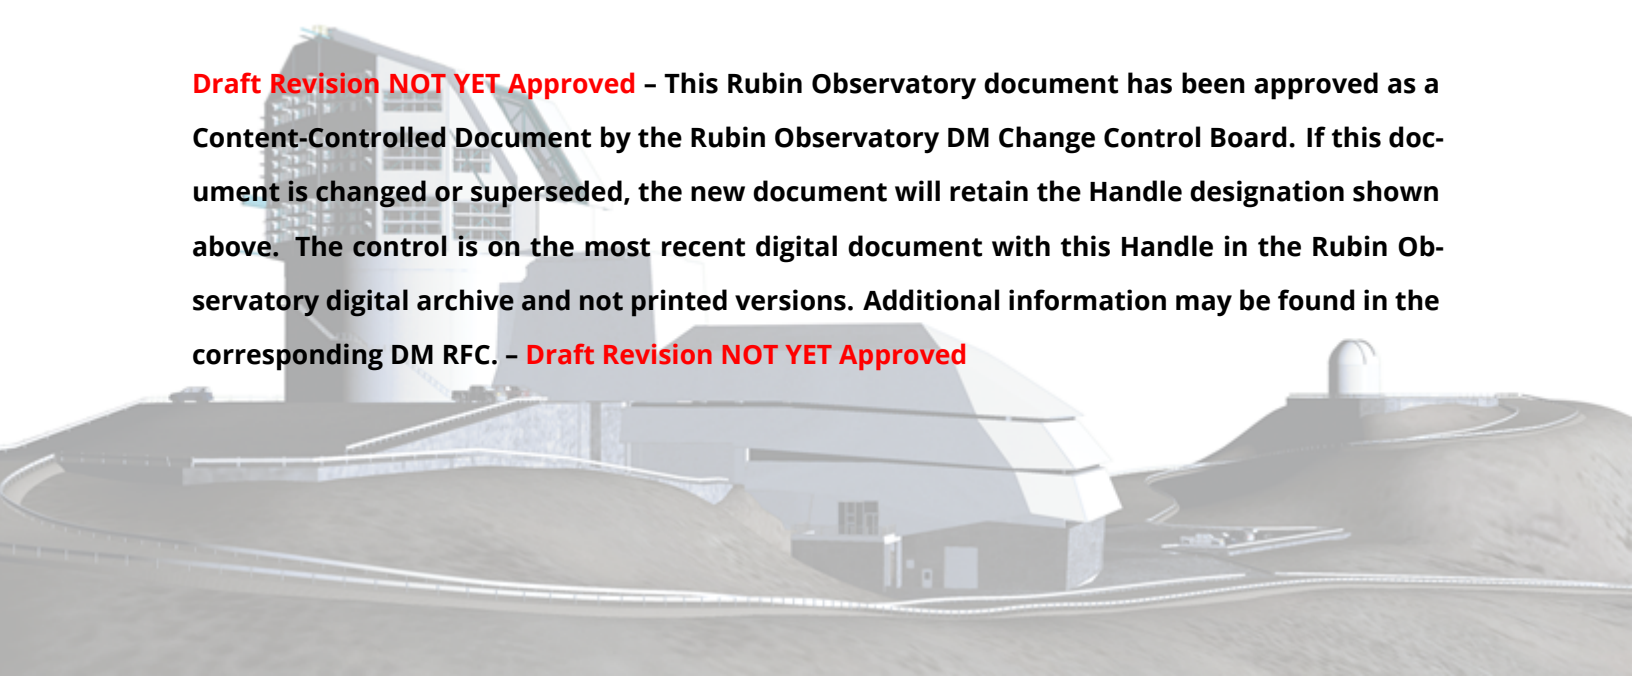
Plans and Policies for LSST Alert Distribution

Eric Bellm, Robert Blum, Melissa Graham, Leanne Guy, Željko Ivezić,
William O'Mullane, Maria Patterson, John Swinbank, and
Beth Willman *for the LSST Project*

LDM-612

Latest Revision: 2020-10-12

Draft Revision NOT YET Approved - This Rubin Observatory document has been approved as a Content-Controlled Document by the Rubin Observatory DM Change Control Board. If this document is changed or superseded, the new document will retain the Handle designation shown above.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 Rubin Observatory digital archive and not printed versions. Additional information may be found in the corresponding DM RFC. - **Draft Revision NOT YET Approved**



Abstract

A major product of the nightly processing of LSST images is a world-public stream of alerts from transient, variable, and moving sources. Science users will access these alerts through community brokers or through a simple filtering service provided by LSST. This document provides a guide to the plans and policies for the alert distribution system to aid science users, broker developers, funding agencies, and LSST Project personnel. It describes the components of the alert distribution system and the data rights required to access specific scientific products. It provides guidelines for organizations developing community brokers and describes the planned capabilities of the LSST simple alert filtering service and Science Platform to aid science users in planning for LSST science.

Change Record

Version	Date	Description	Owner name
	2018-12-26	Initial draft version	Eric Bellm
1.0	2019-01-08	Approved in RFC-556.	Approver: W. O'Mullane
1.1	2019-05-13	Added referece to DMTN-102 RFC-592.	Eric Bellm
1.2	2019-12-27	Revisions supporting broker call for proposals	Eric Bellm
1.3	2020-02-11	Remove extraneous macros	Eric Bellm
1.4	2020-07-09	Update due date; minor updates.	Eric Bellm

Document curator: E. Bellm

Document source location: <https://github.com/lst/lst/LDM-612>

Version from source repository: e4aa087

Contents

1	Introduction	1
2	Components and Capabilities of the LSST Alert Distribution System	3
2.1	Facilities	3
2.1.1	Data Acquisition, Transfer, and Processing	3
2.1.2	The LSST Data Facility and Data Access Centers	3
2.2	Pipelines and Services	3
2.2.1	Alert Production	4
2.2.2	Solar System Processing	5
2.2.3	Alert Distribution	5
2.2.4	The LSST Alert Filtering Service	6
2.2.5	Forced Photometry	6
2.2.6	The LSST Science Platform	7
2.3	Data Products	7
2.3.1	Alert Packets	7
2.3.2	Processed Images	8
2.3.3	Prompt Products Database	8
2.3.4	Data Release Products	8
2.3.5	Alert Database	9
3	Data Rights to Alert Stream Components	10
4	Guidelines for Community Brokers	11
4.1	The Role of Community Brokers	11
4.2	Requirements for Community Brokers	12
4.3	The Broker Selection Panel	13
4.4	Letters of Intent	13
4.5	Call for Proposals	13
4.6	Required Technical Resources	13

4.7	Evaluation of Community Broker Proposals	14
4.7.1	Scientific Value	14
4.7.2	Availability to the World Community	14
4.7.3	Scientific Validity	15
4.7.4	Integration with the Time-Domain Ecosystem	15
4.7.5	Community Adoption	15
4.7.6	Complementarity	15
4.7.7	Existing Agreements	15
4.8	Determining the Number of Community Brokers	16
4.9	Broker Evolution Throughout Operations	16
4.10	Resources for Proposers	16
4.11	Resources for Selected Brokers	17
4.12	Timeline	17
5	Acronyms	19

Plans and Policies for LSST Alert Distribution

1 Introduction

LSST will discover large numbers of astrophysical objects that move across the sky, change in brightness, or appear as transients. Because of their intrinsic temporal variability, full scientific exploitation of many classes of these events requires rapid discovery and additional follow-up observations. Accordingly, LSST's real-time Prompt Processing pipelines will identify such detections using image differencing and report them using world-public *alerts* issued within 60 seconds of the shutter closure after each visit.

These alert packets will contain not only the information about the most recent detection but also a historical lightcurve, cutout images, timeseries features, and other contextual information. A science user will be able to use the contents of a single LSST alert packet to make a decision that is both rapid and informed about whether the event is relevant to their scientific goals. LSST expects to produce up to about ten million of these alerts nightly¹. The resulting *alert stream* will be large—more than 1 TB nightly—but will contain all classes of astrophysical events, from the youngest supernovae to the most distant RR Lyrae to the slowest-moving Trans-Neptunian Objects to the most unexpected new phenomena. The task for the scientist is to identify the small subset of events of interest in the larger alert stream.

To do so, astronomers will rely on third party *community brokers*, software systems that receive the full LSST alert stream and provide additional information to refine the selection of events of interest. Community brokers may crossmatch the LSST stream to multiwavelength catalogs, join LSST alerts with those from other surveys, provide machine-learned classifications of events, and/or offer user filters to winnow the stream. LSST will itself offer an alert filtering service of more limited capacity that will apply user-supplied filters to the alert stream. Individual scientists may receive alerts through one or more of these services.

The large volume of the alert stream and the finite bandwidth from the LSST Data Facility necessitate a proposal process to select community brokers that will receive the full stream.

¹The LSST alert stream will contain essentially all DIASources detected at 5σ in the difference image, including a currently unknown fraction of artifacts. LSST will provide a threshold that may be used to filter transients in the alert stream to 90% completeness and 95% purity at 6σ [LSE-30], so less than 5% of alerts filtered in this manner will be artifacts. By varying the cutoff spuriousness value users may adopt other tradeoffs between completeness and purity as well.

This document outlines the process, criteria, and timeline by which LSST will choose community brokers. To provide context, we also summarize the major features of the LSST Alert Production systems as well as relevant data rights concerns.

Throughout the document, we provide references to the formal LSST requirements and relevant design documents as an aid to the interested reader.

Draft

2 Components and Capabilities of the LSST Alert Distribution System

In this section we provide a high-level overview of the LSST Alert Production process to provide background for science users of LSST alerts. Key numbers for Alert Production are summarized in DMTN-102.

2.1 Facilities

2.1.1 Data Acquisition, Transfer, and Processing

LSST will survey the sky repeatedly using standard or alternate standard visits (currently baselined as 2×15 second exposures in a single band or a single 30 second exposure, respectively). Immediately after the camera reads out, the images will be transferred via fiber networks to the LSST Data Facility (LDF; §2.1.2). All subsequent data processing leading to alerts occurs in the LDF. Alerts will be distributed from the LDF within 60 seconds of the camera readout completing.

2.1.2 The LSST Data Facility and Data Access Centers

LDM-230 describes the concept of operations for the LSST Data Facility. The LDF hosts a variety of computational services, including the Prompt Processing services that run the Alert Production pipelines (§2.2.1) and the Batch Production services that run the annual Data Release Processing (§2.3.4). All Prompt Processing and Alert Production runs at the LDF; Data Release Processing is split between the LDF and the CC-IN2P3 in France. The LDF also hosts a Data Access Center (DAC) that provides access to proprietary data products (§3) through the LSST Science Platform (§2.2.6). The LDF also hosts the alert stream feeds to community brokers (§2.2.3) and the LSST Alert Filtering Service (§2.2.4).

An additional DAC will be located in Chile. Its capabilities are described in LDM-572. It will host a copy of the raw data and processed data products and run a version of the Science Platform.

2.2 Pipelines and Services

2.2.1 Alert Production

LSST Data Management (DM) Alert Production (AP) processes data during the night and generates alerts in close to real-time. It is described in LSE-163, with detailed descriptions of its constituent pipelines and algorithmic components provided in LDM-151. The basic modules include single-frame image processing (e.g., instrument signature removal, photometric and astrometric calibration), difference image analysis (DIA; the subtraction of a coadded template image constructed from the most recent Data Release from the new visit image), and source detection, association, and measurement [LSE-163]. AP will process all standard and alternative standard visits, whether obtained as part of the main wide-fast-deep survey or a Special Program [DMTN-065] and generate alerts within the required 60 seconds after the camera readout completes (DMS-REQ-0004, LSE-61)². Due to the need for Data Release Production-derived templates, Alert Production cannot run at full scale and full fidelity during neither commissioning nor the first year of operations. LSST DM is currently investigating options for Alert Production in year one [DMTN-107].

All sources in a difference image that have a signal-to-noise ratio $SNR > 5$ in positive or negative flux are considered “detected” DIASources, are incorporated into the DIASource catalog, and will cause an alert to be issued³. Additionally some $SNR < 5$ sources will be kept for diagnostic purposes, but will not lead to alerts. Source association and measurement occurs prior to alert generation: every DIASource will have one unique match to a DIAObject (stationary object) or known Solar System Object SSObject (moving object; see Section 2.2.2). The stationary source association algorithms will be probabilistic and incorporate motion models for parallax and proper motion [LDM-151]. If no association is possible a new DIAObject will be created. Source measurements including centroid, fluxes, shapes, and other characterization parameters (Table 1 of LSE-163) are made, and the time-evolving parameters of the DIAObject such as the parallax, mean flux, and periodic/non-periodic light curve features, will be updated to include the new associated DIASource. At this point, the DIASource will be used to create an alert packet (§2.3.1), regardless of whether the DIASource is associated with a new

²AP may also be able to process non-standard visits with longer or shorter exposure times provided that, e.g., the visit image can be successfully PSF-matched and differenced with a template image. Alert generation in crowded fields may produce more than the maximum of 10,000 alerts per visit required to be supported by DM; alert generation in this circumstance is still under study.

³There may be exceptions to this rule, such as the following two examples. (1) Sources with $SNR > 5$ that have a “high probability of being instrumental non-astrophysical artifacts” [LSE-163], potentially as determined by a to-be-developed spuriousness or real/bogus classifier [LDM-151], may not produce alerts. The limit on false positives is set by the requirement that the alert stream be 90% complete at 95% purity for DIASources of $SNR = 6$ LSE-61. (2) Sources with $SNR < 5$ that meet other to-be-determined criteria, such as the likelihood of being a potentially hazardous asteroid, could produce alerts.

DIAObject, a previously-existing DIAObject, or an SSObject.

2.2.2 Solar System Processing

After the completion of a night of observing, LSST Solar System Processing (SSP; LDM-156) will attempt to identify *new* SSObjects using the updated Prompt Products Database (§2.3.3). SSP algorithms form tracklets from pairs of single-apparition DIASources taken during one night, generate tracks between nights, and fit orbits to identify and characterize moving objects. Newly identified moving objects are added to the SSObject catalog. SSP will report candidate moving objects to the Minor Planet Center (MPC), where they will be publicly available; will ingest all previously known or externally identified moving objects into the SSObjects catalog; and will use MPC astrometry in the orbital parameter fits. Users interested in newly-discovered asteroids will need to query the SSObject catalog or the MPC to identify them. The first time a new SSObject is discovered, SSP will associate past DIASources with the new SSObject as far back as is practical from the accuracy of the orbit. However, revised alerts will not be issued. Future LSST detections of the new SSObject will produce alerts attributed to the SSObject.

2.2.3 Alert Distribution

Alert packets (§2.3.1) will be queued for distribution to community brokers (§4) and the LSST alert filtering service (§2.2.4). LSST is prototyping a bulk transport system built on the open-source distributed queue system Apache Kafka⁴, with Apache Avro⁵ used as a binary serialization format [DMTN-093]. Initial tests indicate that this system performs effectively at the required scale [DMTN-028].

Due to the finite bandwidth out of the LDF, only a limited number of selected community brokers can receive the full alert stream. It will not be possible for individual science users to subscribe directly to the full LSST alert stream. Rather, science users may access LSST alerts through a community broker, or through the alert filtering service (§2.2.4) if they have data access rights (§3). A minimum of five full streams is required to be supported [LSE-61]. The eventual number of selected community brokers (§4.8) is still to be determined, as it depends on several quantities including available network capacity, packet size and contents, number

⁴<http://kafka.apache.org/>

⁵<http://avro.apache.org/>

of alerts per visit, and the rate of alert distribution. Current alert packets are estimated to be about 82 KB [DMTN-102]. Options to reduce the packet size are under consideration; these include removing postage stamps, reducing or removing history, or including only a subset of the `DIASource` or `DIAObject` record contents.

2.2.4 The LSST Alert Filtering Service

The LSST alert filtering service provides a means of real-time access to a subset of the alert stream. Using the Science Platform (§2.2.6), users can upload simple filters (filters which only use information within the alert packet) and receive on their own computer a real-time stream of alerts which pass their filter. This capability is expected to be useful for science cases requiring real-time notification that may be too specific to be well-served by community brokers. The alert filtering service is required to support a minimum of 100 simultaneous users; limits on resource usage and bandwidth (support for transmission of up to 20 full-sized alerts per visit per filter is required) will be imposed to maintain performance⁶. If necessary, a TAC-like process may be used to allocate filtering service resources [LSE-163]. The LSST alert filtering service will support a version of the `VOEvent` standard⁷ that is current at the time of LSST operations.

2.2.5 Forced Photometry

Measurements of `DIAObjects` that are below the $\text{SNR} > 5$ threshold are available through “forced photometry”: flux is measured at a previously-known position. Forced photometry is performed on difference images in two ways; neither contributes to alerts, but are available to users through the Science Platform. First, during the real-time difference image analysis, forced photometry is performed on every difference image at the locations of all previously-known `DIAObjects` detected in a past interval to be determined. The resulting measurements are stored in the Prompt Products Database (§2.3.3) and are queryable within 24 hours. Second, at the end of the night, “precovery” forced photometry will be performed for all *new* `DIAObjects` on every difference image from the past 30 days. These results are also stored in the Prompt Products database and will be available within 24 hours.

⁶The requirement on the number of simultaneously connected users and number of passed alerts is largely driven by outbound bandwidth limitations from the DAC at NCSA. We are investigating approaches that would support larger numbers of active filters. Users with data rights that do not require real-time notification can also programmatically query the Prompt Products Database to identify `DIAObjects` or `SSObjects` of interest without using the LSST alert filtering service.

⁷<http://www.ivoa.net/documents/VOEvent/>

Additionally, a service shall be provided to users to obtain full-survey precovery for a limited number of user-specified `DIAObjects`, also within 24 hours (DMS-REQ-0341, LSE-61). This service will be made available through the Science Platform.

2.2.6 The LSST Science Platform

The LSST Science Platform (LSP) provides access to the proprietary data products (e.g., §2.3.2–2.3.5) held in the DACs. The Science Platform is described in full in LSE-319 and LDM-554. It provides three means of accessing data: a web-based “Portal” for visual examination and querying of the LSST images and catalogs; an interactive “Notebook” environment with associated computing allocations for running code close to the data; and an Application Programming Interface (API) for programmatic access to LSST data using Virtual Observatory standards.

2.3 Data Products

2.3.1 Alert Packets

The contents of an `Alert` packet are fully described in Section 3.5.1 of LSE-163. We reproduce the list of included data here:

- I. an ID uniquely identifying the `Alert`
- II. the prompt products database ID
- III. the database record of the `DIASource` that triggered the `Alert`, as well as the `filterName` and `programId` of the corresponding `Visit`
- IV. the entire associated `DIAObject` or `SSObject` record from the Prompt Products Database, which include a variety of variability metrics computed on the updated `DIASource` lightcurve.
- V. Any `DIASource` and `DIAForcedSource` records that exist, and difference image noise estimates where they do not, taken from the previous 12 months.
- VI. postage stamps of the difference image and template at the `DIASource` location

The full list of parameters that are measured and included in the `DIASource` and `DIAObject` records are provided in Tables 1 and 2 of LSE-163, respectively. The only trigger for an `Alert`

is the detection of a source in a difference image with $\text{SNR} > 5$. Thus, objects that are saturated in the visit or template image, or objects that are the same brightness in both the visit and template images will not generate an `Alert`. Imperfect subtractions may create artifacts with $\text{SNR} > 5$; these will generally lead to an `Alert`, subject to overall capacity limits. A machine-learned spuriousness score will be provided for each `DIASource` triggering an `Alert`, so users may filter the stream for greater completeness or purity depending on their scientific needs. The spuriousness classifier is required to achieve 90% completeness and 95% purity at 6σ [LSE-30] at a threshold spuriousness value.

2.3.2 Processed Images

Processed visit images (PVIs) and difference images produced by the prompt processing pipeline will be available in the Science Platform to users with data rights within 24 hours. They are held on disk for 30 days to enable the precovery forced photometry service (§2.2.5). PVIs for older visits can be regenerated on-demand.

2.3.3 Prompt Products Database

The `DIAObject`, `DIASource`, and `SSObject` catalogs created by the Alert Generation pipelines are held in a Prompt Products Database along with the outputs of the Forced Photometry. This database is queryable through any of the aspects of the Science Platform.

2.3.4 Data Release Products

Annual⁸ releases of all the LSST data will include processed and stacked images, catalogs of `Sources`, `ForcedSources`, and `Objects` from measurements on the stacked and individual images, as well as calibration information. It will also include a reprocessing of all images with the latest pipelines and full-survey versions of `DIASource` and `DIAObject` catalogs—i.e., the variability characterization parameters are calculated from the entire survey to date. All of these products are queryable through the Science Platform.

⁸There will be an additional initial data release containing data from the first six months of operations [LSO-011].

2.3.5 Alert Database

All alerts will be stored in their full original form in an archive hosted in the DACs. This archive is expected to be of most use for testing broker filters or for studies that need to emulate real-time identification of transients from the alert packets. Users who simply wish to work with large samples of lightcurves will likely prefer to query the DIAObject and DIASource or Object and ForceSource tables directly.

Draft

3 Data Rights to Alert Stream Components

LSST data rights and data access policies are formally described in the LSST Data Rights Policy document, LDO-13, which is in development at this writing. The interested reader is advised to consult that document for details. We summarize here for convenience those aspects of the policy that pertain to LSST alerts. In the case of any discrepancy between this document and the LSST Data Rights Policy document, the LSST Data Rights Policy document shall take precedence.

LSST alert packets (§2.3.1) are a world-public data product. We use “public” here in the sense that alert packets can be freely shared with anyone, by anyone, anywhere, worldwide. However, LSST does not have the bandwidth to serve the full alert stream directly to individuals; thus, the term “public” means “shareable” and should not be misinterpreted as “freely available.”

Instead, the LSST alert stream will be delivered to a to-be-determined set of community brokers (Section 4). Brokers may share (or not share) the contents of alert packets with whomever they choose, although the broker selection process (§4) will prioritize selection of brokers that will provide world-public access to their services. The contents of the DIAObject, DIASource, DIAForcedSource, and SSObject tables held in the PPDB are also public and may be freely shared, although the ability to access and query the PPDBs held in LSST DACs is restricted to data rights holders.

Access to proprietary LSST data (e.g., recent LSST images or the current LSST data release) is restricted to data rights holders. Brokers run by LSST data rights holders may facilitate access to these data for data rights holders but may not redistribute proprietary data. Proprietary data may be used as part of alert classification; the alert classifications are a derived data product and may be made public.

4 Guidelines for Community Brokers

4.1 The Role of Community Brokers

Community brokers will play a vital role in enabling time-domain science with LSST. By receiving the LSST alert stream, adding value, and redistributing it to the scientific community, brokers facilitate full exploitation of the scientific value of the real-time LSST alerts.

We envision that community brokers, individually or jointly, may provide a variety of functions. These may include, but are not limited to:

- redistributing alert packets to a defined or world-wide community
- filtering alerts for specific science goals or to remove potentially spurious alerts⁹
- cross-correlating LSST alerts with other static catalogs or alert streams
- classifying events scientifically
- providing user interfaces to the data
- coordinating scientific activity among collaborators
- triggering followup observing
- for users with appropriate data rights, facilitating followup queries and/or user-generated processing within the LSST Data Access Center
- managing annotation & citation as followup observations are made
- collecting classification and other information gathered by the scientific community

Some of these functions may be performed by “Target and Observation Managers” [TOMs; 18] that may either be integrated into community brokers directly or connect to annotated and/or filtered streams delivered by a broker. For simplicity, we will here refer to any system that connects directly to the LSST alert stream as a broker.

⁹trading completeness for purity

4.2 Requirements for Community Brokers

Because of the large data volume of the alert stream (several TB per night), finite bandwidth from the LSST datacenter prevents sending a full copy of the alert stream to all interested parties. Additionally, we anticipate that community brokers will require some level of support from the LSST operations team. Accordingly, LSST is conducting an open proposal process in order to select brokers with sufficient technical capability to enable valuable and reliable scientific returns.

We believe the scientific community will be best served by a rich broker ecosystem offering diverse capabilities. Any institution worldwide was allowed to submit a Letter of Intent (§4.4). The LSST SAC decided to invite all institutions who submitted a LOI to submit a full proposal. Demonstration of appropriate technical and personnel resources to support the proposed goals were the minimum requirement to be considered.

Recognizing that different science goals may benefit from different technical approaches, and that proposing institutions may have a range of strengths and experience, the proposal solicitation does not place any explicit functional requirements on proposed brokers. Thus a proposed broker is not *required* to consume the whole stream, to redistribute the full stream, to operate for the full ten-year survey, or to make its products world-public, for example. However, given the small and finite number of brokers to be selected (§4.8) and the evaluation criteria (§4.7), proposals offering to do so will be more likely to be selected. The selection committee will evaluate the proposals holistically, weighing the unique capabilities proposed in order to maximize the scientific output of LSST.

While we expect the evaluation process to favor proposals that receive the full alert stream, brokers may propose to receive only a filtered subset of events (or a subset of the alert packet contents) from LSST. We anticipate that this capability would be provided by the same technology underlying the LSST alert filtering service, potentially with less restrictive requirements than are placed on user-provided filters. Some broker services, particularly those processing only a small subset of alerts, may elect to receive alerts from other community brokers rather than from the LDF directly. Brokers requesting direct access to the alert stream who agree to forward alerts to downstream services will be favored in the selection process.

4.3 The Broker Selection Panel

Broker proposals will be evaluated by a panel designated by the LSST Science Advisory Committee. Appropriate care will be taken to avoid conflicts of interest. The panel will include ex officio representatives from the LSST Project Office, LSST Operations, and LSST Data Management to provide policy and technical guidance.

4.4 Letters of Intent

The broker proposal process will have two stages: an initial open call for Letters of Intent (LOIs) from all interested parties, and a subsequent full proposal call solicited from invited LOI writers. This two-stage process ensures those writing full proposals have a reasonable chance at selection and provides an opportunity for early feedback. An invitation to submit a full proposal may also be valuable to proposers in obtaining relevant funding. No financial support is available from the LSST Project. The LOI call was issued in LDM-682.

LOIs were evaluated using the same criteria as the full proposals (§4.6–4.7), with recognition that the proposed brokers may have been in the conceptual or early design stage at the time of the LOI.

4.5 Call for Proposals

The call for full proposals for LSST brokers was issued in LDM-723. The broker selection panel will use these proposals to select the brokers that will receive alert streams directly from the LSST Data Facility.

4.6 Required Technical Resources

The panel will evaluate whether the proposed broker system is technically feasible and appropriately supported for the proposed goals. At minimum, this requires:

- Large inbound and outbound network bandwidth (the full alert stream is a few TB/night)
- Petabytes of disk capacity

- Databases capable of handling billions of sources
- Compute resources to handle sophisticated classification and filtering tasks in real time at scale
- Appropriate personnel and sufficient effort to develop the required software and to maintain and operate the service
- Institutional support to ensure the longevity and stability of the service and software. By default selected brokers will be expected to plan to operate for the full ten-year survey.
- Discussion of funding sources to support broker development and operations.

4.7 Evaluation of Community Broker Proposals

The broker selection panel will evaluate proposals with a primary goal of maximizing the scientific exploitation of LSST—recognizing that many of LSST’s scientific returns are likely to be unanticipated.

The panel will evaluate broker proposals that demonstrate appropriate technical resources (§4.6) using the following evaluation criteria:

4.7.1 Scientific Value

The panel will consider whether a proposed broker has the potential to add scientific value that serves a large community, enables high-profile science, or provides unique capabilities. In particular, the panel will evaluate the proposed contributions to LSST’s four science pillars¹⁰. The panel will prioritize proposals that require or take advantage of the unique aspects of the LSST alert stream, namely its real-time and world-public nature.

4.7.2 Availability to the World Community

Because community brokers are the means by which scientists without LSST data rights may access the public contents of the alert packets, the panel will prefer proposals that make their

¹⁰These are: Probing Dark Matter and Dark Energy, Taking an Inventory of the Solar System, Exploring the Transient Optical Sky, and Mapping the Milky Way [LPM-17].

products and services world-public. At least one broker that provides world-public access to alerts will be selected.

4.7.3 Scientific Validity

The panel will look for evidence that scientific products, such as photometric classifications, produced by a proposed broker are accurate. Brokers may demonstrate performance in production on precursor surveys or by data challenge. The proposal call will provide sample alerts generated by LSST pipelines from precursor survey data that may be used for testing and evaluation.

4.7.4 Integration with the Time-Domain Ecosystem

For science cases involving real-time followup, the panel will consider the extent to which proposed brokers facilitate integration with followup resources, other surveys and alert streams, other brokers and services, existing communities of observers, and archives.

4.7.5 Community Adoption

The panel will consider whether a proposed broker system has already demonstrated scientific value and community adoption on precursor streams. This evaluation will also assess the range of scientific, geographic, and institutional communities that have utilized the service as well as the publications it enabled.

4.7.6 Complementarity

To maximize scientific returns, the panel will aim to select an ensemble of complementary brokers. Brokers that are more limited in scope or scientific goals may be selected if they provide capabilities not present in other selected brokers or are serving scientific communities that would otherwise be unable to use LSST alerts.

4.7.7 Existing Agreements

The panel will consider the applicability of any existing agreements, if relevant.

4.8 Determining the Number of Community Brokers

Early sizing estimates indicated that only four full streams could be provided from the LSST Data Facility due to network bandwidth limitations. Current capacity projections indicate a larger number may be possible (§2.2.3), but with some uncertainty. A minimum of five full streams is required [LSE-61]. The available capacity will depend on:

- actual network capacity and usage patterns
- details of the alert format
- the achieved purity of the alert stream

Final capacity determination will not be possible until commissioning data is available. The broker selection panel may elect to make guaranteed and provisional selections based on these estimates. A subset of brokers may be selected for integration during the commissioning and early operations phases, with additional brokers joining as the survey reaches steady state.

4.9 Broker Evolution Throughout Operations

Broker performance and usage will be reviewed periodically during LSST Operations at an interval to be determined. Broker selections may be changed or additional new brokers added due to changes in LSST survey or pipeline performance, broker performance or usage, or community science priorities.

4.10 Resources for Proposers

To aid in broker development and testing, LSST is providing pre-packaged alert streams for precursor datasets generated by the LSST Alert Production pipelines and distributed with the expected formats and protocols. These LSST-processed alerts provide the best match to the content and formats currently envisioned for LSST alerts, although the precursor surveys themselves do not closely match the LSST depth, filters, cadence, etc. Details on these sample alerts may be found at <https://www.lsst.org/scientists/alert-brokers>.

Proposers may also wish to take advantage of the public alert archive from the Zwicky Transient Facility¹¹. ZTF alerts are shallower than the sample precursor alerts produced by the LSST pipelines and contain somewhat different information¹². However, ZTF alerts are available in much greater quantities and in many cases have been spectroscopically classified¹³.

Proposers may also or instead utilize simulations or other precursor datasets in their proposals if desired. In any case, proposers should be explicit about what data they use for demonstrations, and explain the extent to which, and the limitations to which, they illustrate the scientific and technical capabilities of their system.

4.11 Resources for Selected Brokers

Selected brokers will be expected to sign a Memorandum of Understanding codifying agreement to respect LSST Data Rights policies (where relevant) and outlining expected interfaces, support, and Service-Level Agreements for both parties. Selected community brokers will receive basic support from the LSST Operations Team to ensure that the alert stream is delivered with the expected form, content, reliability, and timing. Additionally, status information will be available during operations to diagnose processing and data transport problems as they arise.

4.12 Timeline

The calendar-year timeline presented below is the current best estimate based on the LSST project schedule and is subject to revision.

- 2018** • Q4 Plans and Policies for LSST Alert Distribution document issued
- 2019** • Q1 Issue call for Letters of Intent (LOI)
 - Q2 Letters of intent due
 - *June* Broker development workshop (invited LOI writers)
 - Q4 Revise this document. Invite full proposals for community brokers from selected LOIs

¹¹<https://ztf.uw.edu/alerts/public/>

¹²ZTF alerts have 30 days of history rather than one year, and no forced photometry or timeseries features, but do include crossmatches to Gaia and PanSTARRS.

¹³See e.g., the Transient Name Server (<https://wiserep-tns.weizmann.ac.il/>)

- 2020**
 - Q3 Document bulk transport format and interface. Distribute alert stream test software.
 - *December 15* Broker proposals due.
- 2021**
 - Q2 Provisional broker selections made.

Subsequent steps depend on major commissioning milestones and will shift according to the evolving commissioning schedule¹⁴. Alerts are provided during commissioning to enable engineering tests of broker interfaces. Potential data release scenarios are described in LSO-011.

Draft

¹⁴See <https://www.lsst.org/about/project-status>.

5 Acronyms

Acronym	Description
AP	Alerts Production
API	Application Programming Interface
C	Specific programming language (also called ANSI-C)
DAC	Data Access Center
DAPOL	Data Access Policy (Tag for policy statement)
DIA	Difference Image Analysis
DM	Data Management
DMS	Data Management Sub-system
DMTN	DM Technical Note
DRP	Data Release Production
ID	Identifier (Identification)
KB	KiloByte
LDF	LSST Data Facility
LDM	LSST Data Management (handle for controlled documents)
LOI	Letter of Intent
LPM	LSST Project Management (Document Handle)
LSE	LSST Systems Engineering (Document Handle)
LSP	LSST Science Platform
LSST	Large Synoptic Survey Telescope
MOPS	Moving Object Pipeline System
MPC	Minor Planet Centre
NCSA	National Center for Supercomputing Applications
PSF	Point Spread Function
RR	Rate Reduction
SNR	Signal-to-Noise Ratio (also denoted SN and S/N)
TAC	Time Allocation Committee
TB	TeraByte

References

- [1] **[LDM-682]**, Bellm, E., Blum, R., Graham, M., et al., 2019, *Call for Letters of Intent for Community Alert Brokers*, LDM-682, URL <https://ls.st/LDM-682>

- [2] **[LDM-723]**, Bellm, E., Blum, R., Graham, M., et al., 2019, *Call for Proposals for Community Alert Brokers*, LDM-723, URL <https://ls.st/LDM-723>
- [3] **[LDO-13]**, Blum, R., et al., 2019, *LSST Data Policy*, LDO-13, URL <https://ls.st/LDO-13>
- [4] **[LSE-30]**, Claver, C.F., The LSST Systems Engineering Integrated Project Team, 2018, *Observatory System Specifications (OSS)*, LSE-30, URL <https://ls.st/LSE-30>
- [5] **[LSE-61]**, Dubois-Felsmann, G., Jenness, T., 2018, *LSST Data Management Subsystem Requirements*, LSE-61, URL <https://ls.st/LSE-61>
- [6] **[LDM-554]**, Dubois-Felsmann, G., Ciardi, D., Mueller, F., Economou, F., 2018, *Science Platform Requirements*, LDM-554, URL <https://ls.st/LDM-554>
- [7] **[DMTN-065]**, Graham, M., Jurić, M., Lim, K.T., Bellm, E., 2018, *Data Management and LSST Special Programs*, DMTN-065, URL <https://dmtn-065.lsst.io>, LSST Data Management Technical Note
- [8] **[DMTN-102]**, Graham, M.L., Bellm, E.C., Guy, L.P., Dubois-Felsmann, C.T.S.G.P., the DM System Science Team, 2019, *LSST Alerts: Key Numbers*, DMTN-102, URL <https://dmtn-102.lsst.io>, LSST Data Management Technical Note
- [9] **[DMTN-107]**, Graham, M.L., Bellm, E.C., Slater, C.T., the DM System Science Team, 2019, *Options for Alert Production in LSST Operations Year 1*, DMTN-107, URL <https://dmtn-107.lsst.io>, LSST Data Management Technical Note
- [10] **[LPM-17]**, Ivezić, Ž., The LSST Science Collaboration, 2018, *LSST Science Requirements Document*, LPM-17, URL <https://ls.st/LPM-17>
- [11] **[LSE-319]**, Jurić, M., Ciardi, D., Dubois-Felsmann, G., 2017, *LSST Science Platform Vision Document*, LSE-319, URL <https://ls.st/LSE-319>
- [12] **[LSE-163]**, Jurić, M., et al., 2017, *LSST Data Products Definition Document*, LSE-163, URL <https://ls.st/LSE-163>
- [13] **[LDM-156]**, Myers, J., Jones, L., Axelrod, T., 2013, *Moving Object Pipeline System Design*, LDM-156, URL <https://ls.st/LDM-156>
- [14] **[LDM-572]**, O'Mullane, W., Petravick, D., 2017, *Chilean Data Access Center*, LDM-572, URL <https://ls.st/LDM-572>

- [15] **[DMTN-093]**, Patterson, M., Bellm, E., Swinbank, J., 2018, *Design of the LSST Alert Distribution System*, DMTN-093, URL <https://dmtn-093.lsst.io>, LSST Data Management Technical Note
- [16] **[DMTN-028]**, Patterson, M.T., 2018, *Benchmarking a distribution system for LSST alerts*, DMTN-028, URL <https://dmtn-028.lsst.io>, LSST Data Management Technical Note
- [17] **[LDM-230]**, Petravick, D., Butler, M., Gelman, M., 2018, *Concept of Operations for the LSST Data Facility Services*, LDM-230, URL <https://ls.st/LDM-230>
- [18] Street, R.A., Bowman, M., Saunders, E.S., Boroson, T., 2018, In: *Software and Cyber-infrastructure for Astronomy V*, vol. 10707 of Proc. SPIE, 1070711 (arXiv:1806.09557), doi:10.1117/12.2312293, ADS Link
- [19] **[LDM-151]**, Swinbank, J.D., et al., 2017, *Data Management Science Pipelines Design*, LDM-151, URL <https://ls.st/LDM-151>
- [20] **[LSO-011]**, William O'Mullane, L.G., Phil Marshall, 2019, *Release Scenarios for LSST Data*, LSO-011, URL <https://lso-011.lsst.io>