



Vera C. Rubin Observatory
Data Management

Options for Alert Production in LSST Operations Year 1

Melissa L. Graham, Eric C. Bellm, Colin T. Slater, Leanne P. Guy, and
Rubin Observatory Data Management System Science Team

DMTN-107

<https://doi.org/10.71929/rubin/2997867>

Latest Revision: 2020-04-14



Abstract

This document reviews five options for alert production in LSST Operations Year 1 (LOY1), taking into account any implications on LSST formal requirements including up-scopes, down-scopes or explicit violations. The Data Management System Science Team's preferred option for maximizing LSST science is to generate template images from as much of the data from the commissioning and/or science verification phases as possible and use them to run Difference Image Analysis and alert production during LOY1. A proposal to increase the sky area covered during commissioning via a "filler" scheduler program is also presented. As a potential moderate up-scope, this study presents an option to build interim templates on a $\sim$ monthly basis during LOY1, which could increase the accessible sky area by ~ 1000 - 2000 deg^2 per month, and should be reconsidered closer to the start of Operations.

Change Record

Version	Date	Description	Owner name
0	2019-03-28	Internal working document.	Melissa Graham
0.1	2019-04-24	Review.	Leanne Guy
1	2019-05-07	Reviewed by co-authors, ready for circulation.	Melissa Graham
1.1	2019-06-03	Added DMS-REQ-0271,0274 to S.2.2.	Melissa Graham
1.2	2020-04-14	Updated to account for the changed plans for LSST commissioning (Jira ticket DM-22193).	Melissa Graham
1.2.1	2025-10-08	Add DOI	Tim Jenness

Document source location: <https://github.com/lsst-dm/dmtn-107>

Contents

1 Introduction	1
2 The Relevant Requirements, Specifications, and Plans	3
2.1 Alert Production	3
2.2 Alert Packet Contents	3
2.3 Template Generation	4
3 Options for Alert Production in LOY1	6
3.1 Do Nothing	8
3.2 Build Templates from Commissioning Data	10
3.2.1 Optimizing Science with Commissioning-Data Templates	12
3.3 Build Templates with LOY1 Data	15
3.4 Image-Image Differencing	17
3.5 Catalog Differencing	19
4 Summary and Conclusions	20
4.1 The DM System Science Team’s Preferred Option for Alerts in LOY1	20

Options for Alert Production in LSST Operations Year 1

1 Introduction

One of the four science pillars of LSST is time-domain astronomy, which is enabled by alerts on LSST detections of transient, variable, and/or moving objects. Alerts are the only data product that will be immediately available (within 60 seconds of image readout) and publicly shareable (not subject to a proprietary period; LSE-163). The worldwide community is — and has been for a while — actively preparing to process the LSST alert stream and use it to generate groundbreaking scientific results.

Alerts are produced by the LSST Prompt pipelines via Difference Imaging Analysis (DIA), which creates difference images by subtracting a template image from each new image, and then identifies sources in the difference image. Template images are not just any other image of the same field; they are generated by stacking previous images in such a way as to “remove” the flux of any transient or moving objects and “standardize” the flux of variable objects. Using deep stacks as a template also minimizes the amount of noise that it contributes to the difference image (e.g., DMTN-006, DMTN-021). For the LSST, template images will either be generated for a range of airmass, or be subjected to a model correction for airmass, in order to mitigate the effects of differential chromatic refraction (DCR; DMTN-037).

During LOY1 there will not be many pre-existing LSST images from which to build templates¹. The current plans for commissioning would enable template generation over some fraction of the wide-fast-deep (WFD) survey area (LSE-79). During operations, images for template generation will be generated during the annual Data Release (DR) processing, which is scheduled to begin at 0.5 years (DR1), at 1 year (DR2), and on a yearly basis thereafter (LSO-011). It is currently undefined how DIA in Prompt processing — and thus alert production — should proceed in LOY1. The ability of the Data Management System (DMS) to produce alerts during LOY1 depends on the fraction of the WFD survey area covered during LSST commissioning, and the ability to increase that area depends on the processing latency for DR1, which might not finish before the end of LOY1 (LSO-011).

¹Building templates from other instrument’s images is considered fully out of scope and not included in this study

Therefore, the role of this document is to review the formal requirements regarding alert production during LOY1 (§ 2), and to propose and assess all possible approaches to alert production in LOY1 (§ 3), with the goal of maximizing LSST science. A summary of the pros and cons of the various options, as well as the DM System Science Team preferred option for maximizing LSST science is provided in (§ 4).

2 The Relevant Requirements, Specifications, and Plans

The existing formal requirements on alert production, alert packet contents, and template images are described below. Five options for alert production in LOY1, presented in § 3, will each be assessed with respect to these requirements.

2.1 Alert Production

The Science Requirements Document LPM-17 states in Section 3.5 Data Processing and Management Requirements, that the prompt processing *"data products are generated continuously every observing night, including alerts to objects that have changed brightness or position"*. The introductory paragraph of this section also states: *"Detailed requirements on data processing and management will be described in the LSST System Requirements Document. Here, only a rough guidance is provided."* For this reason, we do not consider the LPM-17 statement on nightly alert generation be a hard requirement that alerts be generated during LOY1, prior to the availability of the necessary templates.

The LSST Observatory System Specifications document LSE-30 details formal requirements on the fraction of science visits that experience delayed and failed alert production, but these only apply when alert production is *possible* for a visit (i.e., when a suitable template exists).

sciVisitAlertDelay
sciVisitAlertFailure
OSS-REQ-0112

There are also formal requirements that the DMS derive and supply threshold values for a spuriousness parameter (a real/bogus score for each difference image source), which can be used to filter the alert stream into a subsample of transient and variable objects with a given completeness and purity². Substantial changes to the template images (i.e., in depth or stacking algorithm) may require that the real/bogus algorithm be retrained, and the spuriousness threshold parameters be reevaluated, in order to meet these formal requirements on the DMS.

transSampleSNR
transCompleteness-
Min
transPurityMin
OSS-REQ-0353

2.2 Alert Packet Contents

LSE-61 contains a formal requirement that the alert packets include *at least* a 12 month history (if available), as well as cutouts of the template and difference images. The availability of the 12 month history and the cutouts are also discussed in Section 3.5.1 of LSE-163. The 12 month

DMS-REQ-0274

²See also OSS-REQ-0354 for the required parameters for a subsample of transient and variable objects LSE-30.

history would not be available in LOY1, and the requirement adequately accounts for this, so it would not be violated in LOY1. Alerts issued without corresponding image cutouts being made available, either as part of the alert packet or via an alternative distribution mechanism, would violate this requirement.

LSE-61 contains a formal requirement that the variability characterization parameters in an alert packet be based on data collected during *at least* the past year. In LOY1, releasing variability characterization parameters in the alerts that have been calculated with < 1 year of data would, technically, be in violation of this requirement. However, that is unavoidable no matter which option for alert production in LOY1 is adopted.

diaCharacterization-
Cutoff
DMS-REQ-0319

LSE-61 contains a formal requirement that all DIAObjects shall include identifiers of the nearest three stars and three galaxies from the Data Release Objects catalog, and these DIAObject records will be included in the alert packet. Prior to the generation of DR1 in LOY1 this requirement cannot be met, and the alert packets will not contain associations to LSST DR Objects. As in the above paragraph, this is unavoidable no matter which option for alert production in LOY1 is adopted.

DMS-REQ-0271

DMS-REQ-0274

2.3 Template Generation

There are formal requirements on the depth of a template image to be used in DIA, which are expressed as constraints on the fraction (f) of the noise that a template may contribute to the difference image: $f_{Y1} = 0.4$ in year 1, and $f = 0.2$ thereafter. This requirement specifies that the variance of the difference image (σ_{diff}^2) and the variance of a single science image (σ_{sci}^2) shall be related as: $\sigma_{\text{diff}}^2 \leq (1 + f)\sigma_{\text{sci}}^2$. The variance of the reference (template) image is $\sigma_{\text{ref}}^2 = \frac{1}{N}\sigma_{\text{sci}}^2$, where N is the number of science images combined to create the reference image. Since $\sigma_{\text{diff}}^2 = \sigma_{\text{ref}}^2 + \sigma_{\text{sci}}^2$, we arrive at the expression $N \geq \frac{1}{f}$. During LOY1, $N_{Y1} \geq \frac{1}{f_{Y1}} = 2.5$ and during the remaining years of the survey, $N \geq 5$ images. This is a *theoretical* minimum number of images from which the template should be built, but it is worth noting that this minimum of 5 images for a template matches the minimum number images per year for the least-sampled filter of the WFD main survey, u -band [9]. In the context of this DMTN, we use a baseline of ~ 3 visits to estimate the area of sky that could be covered by templates from, e.g., the commissioning surveys, during LOY1. It is important to note that the training and science validation of the real/bogus algorithm discussed in § 2.1 will determine when alert production will be able to provide thresholds which meet the completeness and purity requirements (section 2.1), and that this is not solely or directly dependent on the template

OSS-REQ-0158
templateNoise-
LevelY1
templateNoise-
LevelY2

image depth.

There are two additional template-related policies which are relevant to this discussion, but which do not pose any potential problems for alert production options in LOY1. The first is that LSE-61 requires that the range of epochs which may contribute to a template image is limited to no more than 1 year. This does not pose a potential problem to LOY1 alerts because creating template images from <1 year's worth of epochs is not in violation of this requirement. The second is that LDM-151 states that TemplateCoadd images should be within a default 2 years of the current CalExp image which is about to undergo DIA (Section 3.2.3 of [15]). This does not pose a problem since any template used in DIA during LOY1 will have been obtained within 2 years.

templateMaxTimes-
pan
DMS-REQ-0280

During LSST Operations the production of template images (TemplateCoadds) is planned to occur during Data Release processing (Section 5, LDM-151); however, some fraction of the WFD main survey area will be imaged during commissioning (§ 3.2). Processing for Data Release 1 (DR1), which will be based on the first six months of data, is expected to begin halfway through LOY1, and DR2, which will be based the first 12 months of data, is expected to begin at the end of LOY1. Processing is not instantaneous and will take several months; DR1 is expected to be released one year after the start of operations, and DR2 towards the end of LOY2 (LSO-011). This is echoed in the Plans and Policies for Alert Generation and Distribution document, which states that *"Due to the need for Data-Release-derived templates, Alert Production cannot run at full scale and full fidelity during commissioning and the first year of operations"* (Section 2.2.1, LDM-612).

The method of template generation might also influence the quantity of alerts that could be produced in LOY1. LDM-151 specifies that in order to use a template image in DIA, it must satisfy one of two conditions in order to mitigate the effects of differential chromatic refraction (DCR). It must either (1) match the position and spatial extent of the new image and have been generated from images with a similar airmass and parallactic angle, *or* (2) have been generated with the inclusion of a model that accounts for DCR (LDM-151). The former would be more restrictive on the area of sky with viable templates for use in LOY1.

3 Options for Alert Production in LOY1

The potential solutions considered herein are:

1. **Do nothing.** Do not generate alerts prior to the release of DR1 (which is expected to be released one year after the start of operations).
2. **Build templates from commissioning and/or science verification data** obtained with the LSST camera by the start of LOY1. Generate alerts if and when a template exists for a field. (At minimum, use templates already built as a part of science verification).
3. **Build templates with LOY1 data** on a, e.g., ~monthly basis. Generate alerts if and when a template exists for a field. This option could be done in addition to the previous option.
4. **Do not build templates. Instead, use image-image differencing** to generate alerts if and when a prior image exists for a field.
5. **Do not build templates. Do not use DIA. Instead, generate alerts for all new point sources**, or cataloged point sources which have changed in brightness by more than some threshold since the last image ("catalog differencing").

These are each discussed in more detail in the following sections. The main aspects of each option to consider are outlined in the following table:

Scope	Does the option require an expansion of scope in terms of software development, computational processing, or full-time equivalent (FTE) hours?
Risks	Does the option impose any risks on the the ability of the observatory or DM systems to deliver LSST science?
Requirements	Does the option violate any written requirements or specifications?
Consistency	Does the option cause the content or format of the template images or alerts in LOY1 to be different, compared to the rest of the survey?
Science	Does the option maximize time-domain science in LOY1?

Regarding the "Science" category, in this document we consider "maximizing time-domain science in LOY1" to mean maximizing the number of alerts *for unique objects* that are produced

with the same content as DM's formal plans (e.g., difference-image fluxes, spuriousness parameters). However, "maximizing science" could also be interpreted as providing as much LSST data for each object as possible, e.g., more filters and/or epochs per unique object, as discussed in § 3.2.1.

3.1 Do Nothing

Do nothing. Do not generate alerts prior to the release of DR1 (one year after the start of operations).

Scope	No expansion of DM scope.
Risks	Potential risks for the DMS.
Requirements	Not fully consistent with the high-level SRD requirement, but does not strictly represent a violation.
Consistency	Does not provide consistent data products.
Science	No science from alerts (or any DIA product) until after DR1.

Scope – No expansion of DM scope. Although this option is not a reduction of DM scope, it is likely be perceived as a as a degradation of the survey by the community and will certainly fall below their expectations.

As an additional note, LSST Education and Public Outreach (EPO) are not expecting to serve alerts to the community prior to Data Release 1 and the beginning of LSST Alert Stream production (LEP-031), and so this option would not have any scope implications for EPO.

Risks – Potential risks for the DMS. There is a risk to DM in missing the opportunity to learn lessons from running alert production in LOY1 and imposes a higher risk on alert production when it finally does start in year 2. There is another risk to the DMS: it is likely that the community will do ad-hoc template generation, DIA and/or alert production with user-generated pipelines based on the processed standard-visit images released within 24 hours in order to get early science and publications. This might cause additional load on the user computational processing resources available through the Science Platform, especially if multiple user groups attempt to compete in this regard (which seems likely).

L1PublicT
LSR-REQ-0104

Requirements – Not consistent with high-level SRD requirements. As described in Section 2.1, the SRD states that it provides *rough guidance*, rather than a strict requirement, that alerts be generated continuously every observing night (LPM-17). While it is unavoidable that there be some areas without template images during LOY1, and thus also without alerts, the option to not generate *any* alerts in LOY1 is clearly not in the spirit of the SRD requirements, and would certainly not meet the science community's expectations.

Consistency – Does not provide consistent data products. To not produce any alerts — or any DIA Prompt products — in LOY1 would be inconsistent with the other 9 years of the LSST Main Survey. Furthermore, this inconsistency might impose a large down-time on the community brokers, which could impede their ability to commission their systems and/or to garner and maintain funding. Note that the alerts released during LSST commissioning are not required to have the same latency, format, contents, and/or distribution mechanism as required in operations, which might inhibit broker commissioning; brokers will likely be looking to alerts in LOY1 to fulfill this need.

Science – No science from alerts (or any DIA product) until after DR1. If LSST is not going to do any difference imaging in LOY1, this would prohibit science that requires light curves with any underlying source (extended or point) subtracted, such as supernovae. Time-domain science with, e.g., variable stars, AGN, and moving objects would likely proceed using direct image source catalogs.

3.2 Build Templates from Commissioning Data

Build templates from commissioning and/or science verification data obtained with the LSST camera³ by the start of LOY1. Generate alerts if and when a template exists for a field. At minimum, use templates already built as a part of science verification.

Scope	Potential minor expansion of DM scope.
Risks	No risk for the observatory or the DMS.
Requirements	Does not violate any formal requirements.
Consistency	Provides somewhat consistent data products.
Science	Enables some science from DIA and alerts in LOY1.

Scope – Potential minor expansion of DM scope. The commissioning and science verification plans include building template images to test DIA and alert generation in the Prompt pipeline. These templates could be used for DIA during LOY1, and that would not constitute an expansion of scope. However, to maximize the depth and areal coverage of available templates at the start of LOY1, DM could build new templates from *as much* of the commissioning data as possible. Building new templates would constitute a minor expansion of DM scope in terms of processing and FTE. Either of these options for using commissioning data might also lead to a small increase in the scope because it introduces the opportunity to optimize the commissioning survey areas for template-building, which would require additional time for study, debate, and decision-making. Preliminary considerations for optimizing science with commissioning-data templates is discussed further in Section 3.2.1.

Risks – No risk for the observatory or the DMS. As with the option in § 3.1 to Do Nothing, if LSST does not make use of *all* commissioning data to generate templates and run DIA then there may be members of the community who run user-generated pipelines to do so. This might cause additional load on the users’ computational resources. This is much less likely to happen than in the scenario in which no alerts are produced.

Requirements – Does not violate any formal requirements. Although there are formal requirements which state that the alert contents must contain *at least* a 12-month history and that the variability characterization parameters be based on data from *at least* the past year, the correct interpretation of these requirements is that this historical data be included *when it ex-*

³I.e., the full LSSTCam and not the single-raft ComCam which precedes it.

ists. Therefore, releasing alerts with <1 year of history during LOY1 should not be construed as a violation of the alert content requirements.

Consistency – Provides somewhat consistent data products. The DIA data products that are enabled by commissioning templates would be *mostly* consistent in terms of contents, format, and quality, but inconsistent in terms of area covered (and the alert stream might not be continuous, as there would be no alerts for visits without existing templates). Prior to DR1 no alert packets would have an association with LSST static-sky objects (i.e., no potential host galaxy/star information), but it is not a formal requirement that such associations be included in the alert packet, and this situation is unavoidable for all options for alert production prior to DR1. One other potential inconsistency is that when all of the images used to generate a template are from a short time window it is impossible to fully “remove” the flux of any transient or moving objects and “standardize” the flux of variable objects; this would not be an issue with DIA data products later in the survey.

Science – Enables some science from DIA and alerts in LOY1. This option would enable some science to be done with Prompt data products in the sky areas covered by templates. For example, if templates are built for $\sim 10\%$ of the WFD main survey area in all filters, the LSST alert stream would provide an average of 1 million alerts per night (i.e., $\sim 10\%$ of the full-survey volume; DMTN-102). Preliminary considerations for optimizing science with commissioning-data templates is discussed further in Section 3.2.1.

3.2.1 Optimizing Science with Commissioning-Data Templates

The plans for LSST commissioning surveys are described in LSE-79. When this Tech Note was written, LSE-79 described Final Science Verification Mini-Survey 1 in two phases. First, a three week survey of $\sim 10\%$ of the LSST WFD main survey area ($\sim 1600 \text{ deg}^2$) with ~ 15 visits per filter per field (~ 1 year equivalent depth), from which template images are generated. Then, after a six-week gap, the same three week survey is executed and data is processed with the Prompt pipeline to perform DIA and generate alerts (which might be distributed with high latency via non-streaming interfaces). If Mini-Survey 1 is executed prior to the start of operations, then multi-band templates for $\sim 1600 \text{ deg}^2$ ($\sim 10\%$ of WFD) would already exist and could be used for alert generation in LOY1. LSE-79 also described Mini-Survey 2 as $\sim 300 \text{ deg}^2$ to a 10-year equivalent depth, and a $\sim 100 \text{ deg}^2$ 20-year Depth Test as part of Early Science Verification. Additional sky area would be observed as a part of the other Key Performance Metrics during Early Science Verification. **The main takeaway is that LSE-79 described that at least $\sim 14\%$ of the WFD main survey area could have six-filter templates at the start of LOY1.**

The ability to generate six-filter templates and distribute alerts during LOY1 over 14% of the WFD area would enable some science, but would not maximize it under this Tech Note's working definition of maximizing the number of alerts for unique objects in LOY1. There are two options for increasing this area – and thus the number of LOY1 alerts – prior to the start of operations: (1) dedicate commissioning time to a template-building survey, or (2) create a template-building program for the observation scheduler which can run as "filler" during commissioning.

(1) The Feasibility of a Dedicated Template-Building Survey Prior to Operations

Significantly increasing the sky area with templates by observing additional areas during the already very busy commissioning phase is probably not be feasible. For example, it might be possible to cover $\sim 50\%$ of the WFD area with 3 images per field in a single filter in ~ 3 photometric nights, which should be sufficient to build template images that meet the noise contribution requirement for LSST year 1 (§ 2.3). However, that is a very idealized minimum time estimate, which assumes 1000 visits per night and does not account for weather or dithering, factors which would $\sim$ double the time estimate. Furthermore, such a survey dedicated to template-building might not serve an independent commissioning-related purpose that would make it worth spending (potentially limited) commissioning time on, and priorities would need to be considered.

OSS-REQ-0158
templateNoise-
LevelY1

(2) A Template-Building “Filler” Program During Commissioning

During commissioning, whenever activities that require images be taken but are agnostic about their location and filter (e.g., testing telescope parameters, or troubleshooting scenarios), and a filter change and/or telescope slew would not cause an unacceptable overhead or risk (at the discretion of the commissioning team), then obtain images that would be beneficial to building templates for LOY1 alert production. An observation scheduler program could be written so that this happens automatically whenever the commissioning team chooses this “Template-Building” mode, and the program could optimize achieving, e.g., >3 images in good conditions in a single-filter across as wide an area as possible. It is feasible that the equivalent of ~ 6 photometric nights’ worth of time out of LSSTCam commissioning might be spent in this mode, and as described above this could lead to fully filled single-band templates for $\gtrsim 50\%$ of the WFD main survey area.

The broader question of what it means to “optimize science”.

For the purposes of this Tech Note we have considered “maximizing science” as being equivalent to “maximizing the number of alerts for unique objects”. Under this assumption, the priority for template-building during commissioning would be to build single-filter templates over as wide an area as possible. The optimal filter debatable: it’s likely to be r or i -band, and the latter might have the advantage of not requiring airmass correction for DCR (Section 5.2.4 of LDM-151). Many surveys have done single-band transient searches in the past with much success (e.g., PTF extragalactic focused on r -band).

However, the science community might instead prefer to have more *information* from LSST for each unique object, and prefer the prioritization of multi-band templates over sky area. Furthermore, the rejection of false positives would be aided by multi-filter alert distribution, especially in the early stages as the real-bogus classifier is building up an adequate training set, which would also help to “maximize science”. A time-average of ~ 1 million alerts per night from $\sim 10\%$ of the WFD area will still be a huge increase for the science community to handle, and follow-up resources will be limited. **Further study – including feedback from the science community – is required in order to determine how to prioritize the observing strategy for template-building.** These studies should consider whether to prioritize building templates in galactic or extragalactic fields; for example, galactic variable-star science can still be done with the direct images (although this is more difficult for very crowded fields), but extragalactic transients in extended host galaxies absolutely require image subtraction for detection and analysis. These studies should also consider the observing strategy for LOY1 if templates will be generated on a, e.g., $\sim$ monthly basis as proposed in § 3.3. Since the first

version of this Tech Note, several of the Science Collaborations have begun to compile their science-based recommendations for templates and alerts during LOY1; the Transients and Variable Stars SC has two AAS research notes based on time critical and non-time critical science, Street et al. [13] and Hambleton et al. [7], and the Solar System SC has a white paper describing their plans for high-impact time-domain science in year one⁴.

⁴Schwamb et al. (2020), http://lsst-sssc.github.io/Files/lsst_templates.pdf

3.3 Build Templates with LOY1 Data

Build templates with LOY1 data on a, e.g., ~monthly basis. Generate alerts if and when a template image exists for a field. This option is an extension of building templates with commissioning data (§ 3.2).

Scope	Moderate expansion of scope.
Risks	No risk for the observatory or the DMS.
Requirements	Does not violate any formal requirements.
Consistency	Provides somewhat consistent data products.
Science	Enables some science from DIA and alerts in LOY1.

Scope – Moderate expansion of scope. Since the current DM processing plans are to build templates once during the yearly data release processing effort, it would be a small expansion of scope in terms of computational processing and FTE to evaluate recently obtained data, and generate and validate new template images, on a ~monthly basis prior to DR1. The additional processing resources to stack ~3 images per field over the new ~1000–2000 square degrees obtained each month would likely fit into existing compute budgets, as would the fractional FTE required to generate and validate templates (as it would be part of DM’s activities during LOY1 anyway). This option might also increase the scope of the observing strategy committee, as decisions about whether to optimize the early LOY1 observations for template-building and alert production would require additional simulations, analysis, and debate. This plan would have to be well communicated so that the science community could, e.g., optimize their plans for follow-up as new areas are added to alert production and the average number of alerts per night increases on a ~monthly basis.

Risks – No risk for the observatory or the DMS, as in § 3.2.

Requirements – Does not violate any formal requirements, as in § 3.2.

Consistency – Provides somewhat consistent data products, in the same way as § 3.2. The simplest implementation of this option would be to generate templates for a given area only one time (i.e., when ~3 images have been obtained), but another is to continually improve the templates as more data is obtained (e.g., to fill in gaps left by dithering). In the latter case, template images that are *updated* during LOY1 may cause a step-function in the light curve

histories of some difference-image sources if, for example, there was some residual transient flux in one of the templates. However, this may happen during the later years of Operations when templates are updated with each DR, so that is not a true inconsistency.

Science – Enables some science from DIA and alerts in LOY1, as in § 3.2 but potentially expanding by up to 2000 deg² per month. If obtaining and building template coadded images during commissioning (§ 3.2 and 3.2.1) and LOY1 is prioritized, it might be possible to achieve ~ 100% sky coverage in at least one filter by mid-LOY1.

3.4 Image-Image Differencing

Do not build templates. Instead, use image-image differencing to generate alerts if and when a prior image exists for a field.

Scope	Expansion of scope.
Risks	Potential risks for the DMS.
Requirements	Violates some formal requirements.
Consistency	Does not provide consistent data products.
Science	Enables some science from DIA and alerts in LOY1.

Scope – Expansion of scope. While performing image-image differencing may be possible with the existing science pipelines, the modifications, testing, and validation necessary to release DIA data products generated with image-image differencing would constitute a major expansion of scope for DM (in terms of both FTE and the computation processing required to, for example, train real/bogus for a diverse range of template quality). Furthermore, this option might only be feasible if the survey adopts fixed fields, such that the image-image overlap area is always a substantial fraction. Without fixed field, images would have to be reconstructed into single-visit “templates” with the same area as each new image. This might be possible with the existing science pipelines or might constitute an expansion of scope for DM.

Potential risks for the DMS. This option might requires development of software unbound by formal requirements and might impose additional loads on the computational resources. Furthermore, as in § 3.1 this option might inspire the community to run user-generated pipelines to build templates for subtraction, which could cause additional load on resources.

Requirements – Violates some formal requirements. Using a same-depth single-image as a template results in significantly noisier difference images (as incidentally demonstrated by DMTN-006), and would be in violation of the requirement that the template may contribute no more than 40% of the noise in the difference image during LOY1. Using single-image templates might also compromise or complicate DM’s ability to fulfill the formal requirement that spuriousness (real/bogus) threshold values must be supplied in order to filter alerts down to subsamples of known completeness and purity. I.e., such threshold values depend on the quality of the template image, and using single-image templates would cause a wide dispersion in template quality.

OSS-REQ-0158
templateNoise-
LevelY1

OSS-REQ-0353
transSampleSNR
transCompleteness-
Min
transPurityMin

Consistency – Does not provide consistent data products. The DIA data products such as difference-image flux (and associated parameters) would be fundamentally different from other years if DIA in LOY1 is not done with a “transient-free” template image. Furthermore, real/bogus training would be more complicated without a consistent set of template images and the purity and completeness levels of the alert stream would not be consistent in LOY1 (in area or over time).

Science – Enables some science from DIA and alerts in LOY1. Although this option could yield the maximum number of alerts (assuming that the image-image DIA data products pass all requirements, e.g., characterizing real/bogus), it does not reliably generate scientifically usable light curves for transient objects embedded in their host galaxies (i.e., for cases where the transient is also in the image subtracted). Image-image differencing may be sufficient for variable star science, even when the difference flux is contaminated by the presence of the object in the earlier image, as it still shows whether a source has changed brightness. However, additional algorithms might be needed to generate useful light curves [e.g., 1]. Given that only 3 images are required to produce a template that will meet LOY1 requirements on the maximum fraction of noise contributed by the template to the difference image (3.2.1), this option is likely to entail a lot of additional work and risk for little scientific gain.

3.5 Catalog Differencing

Do not build templates. Do not use DIA. Generate alerts for all new point sources and cataloged point sources that have changed in brightness by more than some threshold since the last image ("catalog differencing").

Scope	Expansion of scope.
Risks	Potential risks for the DMS.
Requirements	Violates some formal requirements.
Consistency	Does not provide consistent data products.
Science	Enables some science from DIA and alerts in LOY1.

Scope – Expansion of scope. New software would need to be developed to generate and validate alerts from single-visit image source catalogs, which would constitute a potentially major expansion in scope in terms of FTE hours (but would fit well within scope in terms of computational processing).

Risks – Potential risks for the DMS. If DIA is not run in LOY1, the risks are the same as the option to Do Nothing in § 3.1.

Requirements – Violates some formal requirements. As this option is similar to image-image differencing (but done with catalog fluxes instead of images), it violates some formal requirements in the same manner as § 3.4.

Consistency – Does not provide consistent data products. Without DIA, the alert packet contents would be fundamentally different in content in LOY1 (i.e., no difference-image fluxes, no spuriousness parameter, no difference-image stamp).

Science – Enables some science from DIA and alerts in LOY1. As with image-image differencing in § 3.4, catalog light curves are more scientifically useful for variable stars and asteroids than for supernovae and objects in crowded fields, which would limit the scientific impact of these alerts in LOY1 (despite potentially producing more of them).

4 Summary and Conclusions

The option for alert production in LOY1 that would provide the largest potential science impact with *no* expansion in DM scope is to only use the template images that will have already been generated as a part of commissioning and science verification. A very small up-scope option is to re-generate template images from *all* of the commissioning data that exists by the start of LOY1, and generate alerts whenever a template exists. The total sky area and filter coverage of templates that could potentially be built from the commissioning and science verification surveys, and further strategies to maximize LOY1 science by optimizing early surveys for template-building, were discussed in § 3.2.1. A further moderate up-scope option is to continue to generate interim template images throughout LOY1 as new data is obtained. This could add $\sim 1000\text{-}2000\text{ deg}^2$ in template images per month. These options to use commissioning and LOY1 data to generate template images, and to run DIA in LOY1, are also favored because they do not pose a risk to the observatory or the DMS, will not violate any formal requirements, and will provide consistent data products compared to the following 9 years of Operations.

The options to produce alerts *without* proper templates and/or without using DIA, although they might at first seem like an easy solution to maximize the number of alerts in LOY1, would actually require larger expansions of scope and engender more potential risks, and are hence disfavored. The option to not produce alerts in LOY1 avoids all potential scope expansions, but is inconsistent with a high-level requirement set forth in the SRD and community expectations, and is not without its own risks, notably, the missed opportunity to learn lessons about alert production early on, and potentially overloaded computational resources as the science community runs user-generated difference-imaging pipelines, and is hence equally disfavored. These disfavored options furthermore fail to provide consistent data products to the users and, in some cases, also violate formal requirements on template images and/or alert contents.

4.1 The DM System Science Team’s Preferred Option for Alerts in LOY1

Template images should be generated prior to the start of Operations from as much of the LSSTCam commissioning data as possible, and used for DIA and alert production during LOY1. A “Template-Building Filler Program” that schedules observations which would be beneficial to template-building should be prepared and used by the commissioning team whenever their

activities are agnostic to imaging filter and telescope pointing. Further study – including input from the science community – is required to determine the scientific priorities for such a scheduler (e.g., prioritizing area in a single filter vs. less area covered by multiple filters; maximizing the number of alerts for unique objects vs. maximizing the amount of LSST information per object). A moderate up-scope option to build interim templates during LOY1 as new data is acquired should also be attempted by DM on a best-effort basis, as this could increase the sky coverage by $\sim 1000\text{-}2000 \text{ deg}^2$ per month.

References

- [1] Barris, B.J., Tonry, J.L., Novicki, M.C., Wood-Vasey, W.M., 2005, *AJ*, 130, 2272 (arXiv:astro-ph/0507584), doi:10.1086/491583, ADS Link
- [2] **[LDM-612]**, Bellm, E., Blum, R., Graham, M., et al., 2020, *Plans and Policies for LSST Alert Distribution*, Data Management Controlled Document LDM-612, NSF-DOE Vera C. Rubin Observatory, URL <https://ldm-612.lsst.io/>
- [3] **[LSE-79]**, Claver, C., The LSST Commissioning Planning Team, 2017, *System Al&T and Commissioning Plan*, Systems Engineering Controlled Document LSE-79, NSF-DOE Vera C. Rubin Observatory, URL <https://ls.st/LSE-79>
- [4] **[LSE-30]**, Claver, C.F., The LSST Systems Engineering Integrated Project Team, 2018, *Observatory System Specifications (OSS)*, Systems Engineering Controlled Document LSE-30, NSF-DOE Vera C. Rubin Observatory, URL <https://ls.st/LSE-30>
- [5] **[LSE-61]**, Dubois-Felsmann, G., Jenness, T., 2019, *Data Management System (DMS) Requirements*, Systems Engineering Controlled Document LSE-61, NSF-DOE Vera C. Rubin Observatory, URL <https://lse-61.lsst.io/>, doi:10.71929/rubin/2587200
- [6] **[DMTN-102]**, Graham, M.L., Bellm, E., Guy, L., et al., 2024, *LSST Alerts: Key Numbers*, Data Management Technical Note DMTN-102, NSF-DOE Vera C. Rubin Observatory, URL <https://dmtn-102.lsst.io/>, doi:10.71929/rubin/2997858
- [7] Hambleton, K., Bianco, F., Clementini, G., et al., 2020, *Research Notes of the American Astronomical Society*, 4, 40, doi:10.3847/2515-5172/ab8129, ADS Link
- [8] **[LPM-17]**, Ivezić, Ž., The LSST Science Collaboration, 2018, *LSST Science Requirements Document*, Project Controlled Document LPM-17, NSF-DOE Vera C. Rubin Observatory, URL <https://ls.st/LPM-17>
- [9] Ivezić, Ž., Kahn, S.M., Tyson, J.A., et al., 2019, *ApJ*, 873, 111 (arXiv:0805.2366), doi:10.3847/1538-4357/ab042c, ADS Link
- [10] **[LSE-163]**, Jurić, M., Axelrod, T.S., Becker, A.C., et al., 2023, *Data Products Definition Document*, Systems Engineering Controlled Document LSE-163, NSF-DOE Vera C. Rubin Observatory, URL <https://lse-163.lsst.io/>, doi:10.71929/rubin/2587118
- [11] **[DMTN-021]**, Reiss, D.J., Lupton, R.H., 2016, *Implementation of Image Difference Decorrelation*, Data Management Technical Note DMTN-021, NSF-DOE Vera C. Rubin Observatory, URL <https://dmtn-021.lsst.io/>, doi:10.71929/rubin/2586490

- [12] **[DMTN-006]**, Slater, C., Jurić, M., Ivezić, Ž., Jones, L., 2016, *False Positive Rates in the LSST Image Differencing Pipeline*, Data Management Technical Note DMTN-006, NSF-DOE Vera C. Rubin Observatory, URL <https://dmtn-006.lsst.io/>
- [13] Street, R.A., Bianco, F.B., Bonito, R., et al., 2020, Research Notes of the American Astronomical Society, 4, 41, doi:10.3847/2515-5172/ab812a, ADS Link
- [14] **[DMTN-037]**, Sullivan, I., 2018, *DCR-matched template generation*, Data Management Technical Note DMTN-037, NSF-DOE Vera C. Rubin Observatory, URL <https://dmtn-037.lsst.io/>
- [15] **[LDM-151]**, Swinbank, J.D., Axelrod, T.S., Becker, A.C., et al., 2020, *Data Management Science Pipelines Design*, Data Management Controlled Document LDM-151, NSF-DOE Vera C. Rubin Observatory, URL <https://ldm-151.lsst.io/>, doi:10.71929/rubin/2587108
- [16] **[LEP-031]**, Team, T.L.E., 2018, *LSST EPO Design*, Education and Public Outreach Controlled Document LEP-031, NSF-DOE Vera C. Rubin Observatory, URL <https://ls.st/LEP-031>