



LARGE SYNOPTIC SURVEY TELESCOPE

Large Synoptic Survey Telescope (LSST)
**LDM-503-3 (Alert Generation) Test
Report**

Eric C. Bellm, John D. Swinbank

DMTR-53

Latest Revision: 2018-01-12

Abstract

This is the test report for LDM-503-3 (Alert Generation), an LSST DM level 2 milestone pertaining to the LSST Alert Production System.

Change Record

Version	Date	Description	Owner name
1.0	2018-01-11	Initial version.	Bellm, Swinbank
1.1	2018-01-12	Released via RFC-429	T. Jenness

Document curator: Eric C. Bellm

Document source location: <https://github.com/lstt-dm/DMTR-53>

Version from source repository: 61f599d



Contents

1 Introduction	1
1.1 Objectives	1
1.2 Scope	1
1.3 System Overview	1
1.4 Applicable Documents	2
1.5 References	2
1.6 Document Overview	2
2 Test Configuration	3
2.1 Documents	3
2.2 Hardware	3
2.3 Software	3
2.4 Input Data	3
3 Personnel	4
4 Overview of the Test Results	5
4.1 Summary Table	5
4.2 Overall Assessment	5
4.3 Recommended Improvements	6
5 Detailed Test Results	6
5.1 AG-00-00	6
5.2 AG-00-05	7
5.3 AG-00-10	8
5.4 AG-00-15	9
5.5 AG-00-20	9
5.6 AG-00-25	10
6 Deviations from test cases/procedures	11
6.1 AG-00-00	11



6.2 AG-00-05	11
6.3 AG-00-10	11
6.4 AG-00-15	11
6.5 AG-00-20	11
6.6 AG-00-25	12

LDM-503-3 (Alert Generation) Test Report

1 Introduction

1.1 Objectives

This document describes the results of tests carried out in late (calendar) 2017 on the LSST Alert Production System in order to assess progress against the LSST DM level 2 milestone LDM-503-3. We report on the success or failure of applicable test cases and assess the state of the software and services tested.

1.2 Scope

The overall test plan for the LSST Data Management system is described in LDM-503. This document specifically refers to the late (calendar) 2017 milestone LDM-503-3, which tests the LSST Alert Production System. The overall LSST Alert Production System test specification is defined in LDM-533. The test plan for LDM-503-3 involves the execution of the entire AP-00 (Small Scale Alert Generation Processing) specification, including the following test cases:

AG-00-00 Installation of the Alert Generation science payload

AG-00-05 Alert Generation Produces Required Data Products

AG-00-10 Scientific Verification of Processed Visit Images

AG-00-15 Scientific Verification of Difference Images

AG-00-20 Scientific Verification of DIASource Catalog

AG-00-25 Scientific Verification of DIAObject Catalog

1.3 System Overview

The LSST Alert Production System is that part of the LSST Data Management system which will be responsible for nightly data processing during LSST operations. The most prominent example of such processing is the generation of alerts from LSST difference images. However, the LSST Alert Production System is also responsible for the distribution and filtering of alerts,

execution of precovery and forced photometry measurements, execution of the MOPS payload, and for the Level 1 Quality Control Service (LDM-148). The LDM-503-3 milestone focuses only on the alert generation part of the system.

Note that we may broadly think of the LSST Alert Production System as consisting of two independent parts: the Prompt and Offline Processing Services, which provide scheduling and workflow services, and the Science Payloads, which contain the algorithmic content. The LDM-503-3 milestone exercises only the science payloads.

1.4 Applicable Documents

- LDM-294 LSST DM Project Management Plan
- LDM-503 DM Test Plan
- LDM-533 LSST Alert Production System Test Specification

1.5 References

- [1] **[LDM-533]**, Bellm, E.C., 2017, *Level 1 System Software Test Specification*, LDM-533, URL <https://ls.st/LDM-533>
- [2] **[DMTR-51]**, Bosch, J., Chiang, H.F., Gower, M., et al., 2017, *LDM-503-2 (HSC Reprocessing) Test Report*, DMTR-51, URL <https://ls.st/DMTR-51>
- [3] **[LDM-148]**, Lim, K.T., Bosch, J., Dubois-Felsmann, G., et al., 2017, *Data Management System Design*, LDM-148, URL <https://ls.st/LDM-148>
- [4] **[LDM-503]**, O'Mullane, W., Jurić, M., Economou, F., 2017, *Data Management Test Plan*, LDM-503, URL <https://ls.st/LDM-503>
- [5] **[LDM-294]**, O'Mullane, W., Swinbank, J., Jurić, M., DMLT, 2017, *Data Management Organization and Management*, LDM-294, URL <https://ls.st/LDM-294>

1.6 Document Overview

Section 2 of this document provides details of the LSST Alert Production System baseline used for this test, including relevant hardware and software configurations. Section 3 lists the in-

dividuals involved in performing the tests. Section 4 provides an overview of the test results, while Section 5 provides more detailed results from each individual test case.

2 Test Configuration

2.1 Documents

This test report refers to the execution of tests AG-00-00 through AG-00-25 in LDM-533 version 1.1.

2.2 Hardware

All tests were executed on systems in the LSST Data Facility.

Software installation (AG-00-00) and scientific analysis work (AG-00-05 through DRP-00-25) were carried out on `lsst-dev01.ncsa.illinois.edu`. At time of text execution, this was a Dell PowerEdge R730 with 24 physical Intel Xeon E5-2690v3 CPU cores at 2.60 GHz and 256 GB of RAM.

Bulk data processing (part of AG-00-05) was carried out on the LSST Verification Cluster (VC). At the time of test execution, the VC provided 48 Dell C6320 nodes, each with 24 physical Intel Xeon E5-2680v3 CPU cores at 2.50 GHz and 128 GB of RAM.

2.3 Software

All systems used for testing — including both `lsst-dev01` and the VC nodes — were running CentOS Linux release 7.4.1708. The `devtoolset-6`¹ toolchain, including GCC² version 6.3.1, was enabled throughout these tests.

2.4 Input Data

Input data for all tests was based on the DeCAM “HiTS” dataset, as described in LDM-533.

¹https://access.redhat.com/documentation/en-us/red_hat_developer_toolset/6/html/6.0_release_notes/

²<https://gcc.gnu.org/>



3 Personnel

All test cases were executed by Eric Bellm (University of Washington).

4 Overview of the Test Results

4.1 Summary Table

TEST CASE ID	PASS/FAIL	COMMENTS
AG-00-00	Pass	Refer to §5.1.
AG-00-05	Pass	One CCD improperly processed, some command executions incomplete due to inappropriate cluster configuration. Test goals substantially met. Refer to §5.2.
AG-00-10	Partial Pass	Some improvements in instrument signature removal needed; §5.3
AG-00-15	Fail	Missing provenance information for difference images. Refer to §5.4.
AG-00-20	Pass	Many required DIASource catalog columns are not implemented. However, this is expected for the current stage of development. Refer to §5.5.
AG-00-25	Pass	Many required DIAObject catalog columns are not implemented. However, this is expected for the current stage of development. Refer to §5.6.

4.2 Overall Assessment

- Installation of the Alert Production pipelines was accomplished without difficulty.
- Execution of the Alert Production pipelines was generally successful. A minor subset of the expected products were not generated due to misconfiguration of the batch processing system, but this is not the mode of execution planned for operations.
- The resulting Processed Visit Images are appropriately processed and masked, in general, but some improvements in ISR are necessary. In particular, crosstalk corrections are not applied, and some masks are missing.
- Required provenance information for difference images are absent, and should be corrected once SuperTask and the new Butler are finalized.
- The DIASource catalogs only contain a subset of the many attributes required (see Table 1).

- The DIAObject catalogs at present only implement raw source association; all “lightcurve” features and related summary/aggregate information is missing.

4.3 Recommended Improvements

- The Alert Production pipelines should be integrated into `lsst_distrib` and high-level documentation provided through `pipelines.lsst.io`.
- Crosstalk corrections should be integrated into the `ap_pipe` pipeline.
- Provenance information for difference images should be provided.
- DIASource and DIAObject catalogs should be enhanced to contain all required attributes.
- DMS-REQ-0270 is untestable as written. Further clarification of the criteria by which sub-threshold DIASources are selected to be stored should be defined.
- DMS-REQ-0331 should be clarified or enhanced—which models should have goodness of fit stored? Which color-color diagrams? Are there other quantities of interested that should be required?

5 Detailed Test Results

5.1 AG-00-00

The Alert Generation Science Payload is available from documented channels: PASSED.

We note that general-user documentation was in review at the time of this test (DM-11592).

The Alert Generation science payload can be installed on LSST Data Facility-managed systems: PASSED.

The string `0k` was returned when executing

```
$ python bin/compare expected/Linux64/detected-sources.txt
```

after running the `demo.sh` on both the cluster head node (`lsst-dev01`) and on an example compute node.

One pytest error was encountered by ap_verify:

```
----- MeasureRuntimeTestSuite.test_valid -----  
  
self = <test_profiling.MeasureRuntimeTestSuite testMethod=test_valid>  
  
def test_valid(self):  
    """Verify that timing information can be recovered.  
    """  
    meas = measure_runtime(self.task.getFullMetadata(), task_name='isr', metric_name='ip_isr.IsrTime')  
    self.assertIsInstance(meas, Measurement)  
    self.assertEqual(meas.metric_name, lsst.verify.Name(metric='ip_isr.IsrTime'))  
> self.assertGreater(meas.quantity, 0.0 * u.second)  
E   AssertionError: <Quantity 0.0 s> not greater than <Quantity 0.0 s>  
  
tests/test_profiling.py:54: AssertionError  
===== 29 tests deselected =====
```

This is a unit test limitation already reported in DM-12848 and does not affect the functionality of the package.

5.2 AG-00-05

Execution of the Alert Generation Pipelines produces the expected data products: PASSED

We performed a Butler get for calexp, deepDiff_differenceExp, and deepDiff_diaSrc dataset types using dataIds constructed all visits and CCDs listed in the appendix of LDM-533.

We found several dataIds where products were not generated:

- For visits 411321, 411371, and 411422, CCDs 48–60 and 62 were missing all three of calexp, deepDiff_differenceExp, and deepDiff_diaSrc products. This failure was due to the time limit being reached on the slurm execution script (which looped sequentially over CCDs for each visit).
- For visit 411657, CCD 47 only was missing all three products. Inspecting the log revealed an error:

```
File "/software/lsstsw/stack3_20171023/stack/miniconda3-4.3.21-10a4fa6/
Linux64/meas_algorithms/14.0-7-g23fdb95+15/python/lsst/meas/algorithms/
measureApCorr.py", line 251, in run
    (name, len(subset2), self.config.minDegreesOfFreedom+1))
RuntimeError: Unable to measure aperture correction for required algorithm
'base_PsfFlux': only 0 sources, but require at least 2.
```

We note the need for increased SLURM timelimits in future executions of this test case. However, this is not regarded as fatal to the overall execution of this test case.

The exception raised for visit 411657, CCD 47 was due to a failure to select sufficient stars for PSF modeling. The root cause of this issue is under investigation on DM-13136. It is not regarded as critical to the successful execution of this test case.

We inspected the DIAObject table by opening the l1db/association.db sqlite database produced by ap_pipe with sqlite3. Issuing the query `select * from dia_objects limit 10;` confirmed that the database was non-empty.

5.3 AG-00-10

Generated PVI's include science, mask, and variance images; background model; zeropoint; PSF, and WCS: PASSED

For all of the generated PVI's, we used Butler calls to check for the existence image, mask, and variance planes; zeropoints, PSFs, and WCS. The background models are persisted separately, and were checked with appropriate Butler calls as well.

All dataids where data was generated in AG-00-05 had all of the required components.

Masks and background are appropriately applied: PARTIALLY PASSED

We chose two visits and five CCDs at random: visit 411673, CCDs 15, 32, 45, 55, and 57, and visit 410915, CCDs 3, 8, 23, 41, and 44.

We used the Butler and DS9 to examine these PVI's.

We note that crosstalk correction for DECam was not enabled at the time of processing, pending the implementation of ticket DM-10299.

We note that some evidently bad columns were left unmasked. This has been ticketed for investigation as DM-10381, which will be scheduled with upcoming development work. These bad columns were regarded as falling outside the acceptable levels for this test.

5.4 AG-00-15

Difference images include provenance information relating to the identity of the input exposures and the PSF matching kernel: FAILED

No direct provenance information about the identity of the calexp and template exposures is available through current Butler interfaces. While the input calexp can be reconstructed by matching datalds from the processed image repository, determining the template applied requires a new call to `getCoaddAsTemplateTask`, with concomitant restrictions on stack versioning, template repository association, etc. We have created a new ticket (DM-13085) to improve this handling.

For each difference image, we used Butler calls to confirm that the PSF matching kernel was stored. No failures were reported from the generated images.

Masks are correctly propagated from the input image: PASSED

We used the visits and CCDs from AG-00-10: visit 411673, CCDs 15, 32, 45, 55, and 57, and visit 410915, CCDs 3, 8, 23, 41, and 44.

We used the Butler and DS9 to compare the PVIs to the resulting differences. The difference image masks were often modified relative to the PVIs, making direct comparison more difficult. However, masks were seen to be propagated, with the notable exception of cosmic ray masks. This behaviour was identified in concurrent DRP milestone testing (DMTR-51), where it was not regarded as fatal to successful execution; it has subsequently been resolved (DM-9953).

5.5 AG-00-20

Measurements are Presented in Flux Units: PASSED

For all of the generated DIASource catalogs, we used Butler calls to check for the existence of `PsfFlux` or `ApFlux` entries. No failures were reported from the generated catalogs.

Each DIASource record contains required attributes: PASSED

DMS-REQ-0269 requires many attributes, which are summarized in Table 1. We determined which attributes were present by manually examining a randomly selected DIASource record retrieved from the Butler.

Many of the attributes required by DMS-REQ-0269 were not present. However, this is expected for the current stage of development. We verified that all attributes which are expected to be generated by the codebase at present were properly present. Those attributes which are missing will be added by future development work.

Faint DIASources satisfying additional criteria are stored: N/A

This test was not performed, due to inadequate requirements specification. Refer to §6.5.

Relevant derived quantities are provided in pre-computed columns: PASSED

Examination of the DIASource record shows two entries, `ip_diffim_PsfDipoleFlux_chi2dof` and `ip_diffim_DipoleFit_chi2dof`, which provide goodness of fit measures as required. Color-color diagrams are not relevant for single-visit DIASource catalogs. The requirement does not specify an exhaustive list of relevant derived quantities and hence cannot be fully tested as written.

5.6 AG-00-25

DIAObjects are recorded with unique identifiers: PASSED

This was verified by comparing the output of `select count(distinct(id)) from dia_objects;` and `select count(id) from dia_objects;` on the Level 1 database. Both returned 546688 objects.

Measurements are Presented in Flux Units: N/A

This test was not performed, as no relevant quantities were stored in the DIAObject table. Refer to §6.6.

Each DIAObject record contains contains an appropriate set of summary attributes: PASSED

The attributes linking to the Object catalog are out of scope for this test, as they require joint

processing with DRP.

The only DIAObject property implemented at the time of test execution was a count of the number of DIASources matched to it. However, all values in the column were 1, which is not expected for a real astrophysical dataset. This was identified as due to a bug whereby association was performed correctly but the metadata was not being appropriately updated. This issue was fixed in DM-13052.

Relevant derived quantities are provided in pre-computed columns: N/A

This test was not performed, as no relevant quantities were stored in the DIAObject table. Refer to §6.6.

6 Deviations from test cases/procedures

6.1 AG-00-00

None.

6.2 AG-00-05

None.

6.3 AG-00-10

None.

6.4 AG-00-15

None.

6.5 AG-00-20

At time of carrying out the test, no “additional criteria” per DMS-REQ-0270 had been specified, so no sub-threshold DIASources were stored. No test was therefore executed to check for the existence of these DIASources. A new ticket, DM-12967, has been created to address the inadequacy of current requirements.

6.6 AG-00-25

No measurements (or aggregated measurements, such as means) were contained in the DIAObject table as implemented at the time of the test. It was therefore meaningless to test in which units they were stored.

Similarly, no pre-computed derived quantities were stored as of this test. However, we note that no such quantities are well specified by current requirements documentation.

Attribute	Pass/Fail	Comment
the identity of the Difference Exposure from which it was derived	FAIL	Only implicitly available by reconstructing dataId
the identity of the associated SSOBJect, if any	FAIL	Out of scope for this test; but no column is present
the identity of the parent Source from which this DIASource has been deblended, if any	PASS	parent column is present
epoch of the observation	FAIL	not present in the catalog (requires denormalization from higher-level metadata)
focal plane position centroid and error (pixel)	PASS	base_SdssCentroid_x, base_SdssCentroid_y, base_SdssCentroid_xSigma, & base_SdssCentroid_ySigma
sky position and associated error (radec)	FAIL	coord_ra and coord_dec are present, but no error is reported
SNR of the detection	FAIL	no SNR is reported
calibrated PS flux and associated error	PASS	base_PsfFlux_flux and base_PsfFlux_fluxSigma
likelihood of the observed data given the PS model	FAIL	not present in the catalog
calibrated aperture flux and associated error	PASS	several values of base_CircularApertureFlux_* are present with errors
calibrated flux and associated error for a trailed source model	FAIL	not present in the catalog
length and angle of the trail	FAIL	not present in the catalog
flux and associated parameters for a dipole model	PASS	several dipole flux values are present
parameters of an adaptive shape measurement and associated error	PASS	base_SdssShape_*
a measure of source extendedness	FAIL	not present in the catalog
the estimated background at the position of the object in the template image with associated uncertainty	FAIL	not present in the catalog
a measure of spuriousness	FAIL	not present in the catalog
flags indicating problems encountered while computing the aforementioned attributes	PASS	many flags are present
calibrated flux and associated error for the DIASource as measured on the Visit image	FAIL	not present in the catalog

TABLE 1: DIASource attributes, per AG-00-20.