



Vera C. Rubin Observatory
Software Test Report

LDM-GEN3: Gen 3 Butler Acceptance Testing Test Plan and Report

Jeffrey Carlin

DMTR-271

Latest Revision: 2022-07-08

DRAFT



Abstract

This is the test plan and report for **Gen 3 Butler Acceptance Testing** (LDM-GEN3), an LSST milestone pertaining to the Data Management Subsystem.

This document is based on content automatically extracted from the Jira test database on 2022-07-08 . The most recent change to the document repository was on 2022-07-12.

Draft

Change Record

Version	Date	Description	Owner name
	2020-10-30	First draft	Robert Gruendl
	2021-01-28	Include new test cycle for LDM-556 requirements	Leanne Guy
	2022-06-07	Test campaign LW-P77 completed and results approved. DM-27337	Jeff Carlin

Document curator: Jeff Carlin

Document source location: <https://github.com/lstt-dm/DMTR-271>

Version from source repository: c7dd5c5

Draft

Contents

1 Introduction	1
1.1 Objectives	1
1.2 System Overview	1
1.3 Document Overview	2
1.4 References	2
2 Test Plan Details	4
2.1 Data Collection	4
2.2 Verification Environment	4
2.3 Related Documentation	4
2.4 PMCS Activity	4
3 Personnel	5
4 Test Campaign Overview	7
4.1 Summary	7
4.2 Overall Assessment	9
4.3 Recommended Improvements	9
5 Detailed Test Results	10
5.1 Test Cycle LVV-C160	10
5.1.1 Software Version/Baseline	10
5.1.2 Configuration	10
5.1.3 Test Cases in LVV-C160 Test Cycle	10
5.1.3.1 LVV-T2264 - Butler Gen3 maturity sufficient to support future pipeline development.	10
5.1.3.2 LVV-T1984 - Demonstrate documentation/examples of Gen3 usage and cp_pipe equivalent.	11
5.1.3.3 LVV-T1982 - Run a pipeline on a single node using pipetask.	12
5.1.3.4 LVV-T1987 - Run Calibration Products Processing (CPP)	13

5.1.3.5	LVV-T1983 - Mini RC2 processing capability	15
5.2	Test Cycle LVV-C162	17
5.2.1	Software Version/Baseline	17
5.2.2	Configuration	17
5.2.3	Test Cases in LVV-C162 Test Cycle	17
5.2.3.1	LVV-T1985 - Verify daf_butler raw data ingest	17
5.3	Test Cycle LVV-C190	24
5.3.1	Software Version/Baseline	24
5.3.2	Configuration	24
5.3.3	Test Cases in LVV-C190 Test Cycle	24
5.3.3.1	LVV-T2503 - Verify Outouts from Test Processing Runs	24
5.3.3.2	LVV-T2502 - Verify Outputs from Science Platform	25
5.3.3.3	LVV-T2501 - Verify Outputs from Data Release Production	26
5.3.3.4	LVV-T2499 - Verify Consistent Output Interface	27
5.3.3.5	LVV-T2498 - Verify Writing FITS tables	28
5.3.3.6	LVV-T2497 - Verify Writing FITS images	28
5.3.3.7	LVV-T2496 - Verify filename invariance	29
5.3.3.8	LVV-T2495 - Verify Combining composite datasets for export	30
5.3.3.9	LVV-T2494 - Verify Strong exception guarantee	31
5.3.3.10	LVV-T2493 - Verify No clobber	32
5.3.3.11	LVV-T2492 - Verify Blocked write operation	33
5.3.3.12	LVV-T2491 - Verify Creation of new DatasetTypes	34
5.3.3.13	LVV-T2488 - Verify access outputs from test processing runs	35
5.3.3.14	LVV-T2487 - Verify Accessing official Data Releases	36
5.3.3.15	LVV-T2486 - Verify Consistent input interface	37
5.3.3.16	LVV-T2485 - Verify Local caching of remote resources	39
5.3.3.17	LVV-T2484 - Verify Local proxy	40
5.3.3.18	LVV-T2483 - Verify Failure on missing input file	41
5.3.3.19	LVV-T2482 - Verify Enabling PipelineTasks to execute	42
5.3.3.20	LVV-T2481 - Verify third party datasets	43

5.3.3.21	LVV-T2480 - Verify Item from Composite Datasets	45
5.3.3.22	LVV-T2479 - Verify Parameterized subset of a Dataset	46
5.3.3.23	LVV-T2478 - Verify I/O using cloud storage	47
5.3.3.24	LVV-T2477 - Verify I/O using distributed file system	48
5.3.3.25	LVV-T2476 - Verify Format Plugability	49
5.3.3.26	LVV-T2474 - Verify Data Discovery for Data Release Production	51
5.3.3.27	LVV-T2475 - Verify Data discovery for test processing runs . . .	52
5.3.3.28	LVV-T2473 - Verify Consistent discovery interface	52
5.3.3.29	LVV-T2472 - Verify Introspection for DatasetExpressions	53
5.3.3.30	LVV-T2471 - Verify Filter by non-DatasetRef Database Entries .	54
5.3.3.31	LVV-T2470 - Verify Dataset overrides	55
5.3.3.32	LVV-T2469 - Verify Multiple parallel input Collections	56
5.3.3.33	LVV-T2468 - Verify Multiple chained input Collections	57
5.3.3.34	LVV-T2466 - Verify enable complete pipeline specification . . .	58
5.3.3.35	LVV-T2467 - Verify DataUnit lookup: processing driven	59
5.3.3.36	LVV-T2464 - Verify multiple simultaneous sky definitions	59
5.3.3.37	LVV-T2465 - Verify pipeline execution in multiple contexts . . .	61
5.3.3.38	LVV-T2461 - Verify Collection Layering: Science Platform	62
5.3.3.39	LVV-T2463 - Verify enabling of different execution environments	63
5.3.3.40	LVV-T2462 - Verify QuantumGraph algorithm	64
5.3.3.41	LVV-T2460 - Verify generating a DAG	65
5.3.3.42	LVV-T2457 - Verify butler instantiation	66
5.3.3.43	LVV-T2456 - Verify execution logging	66
5.3.3.44	LVV-T2455 - Verify pipeline interface available as Python API . .	67
5.3.3.45	LVV-T2454 - Verify pre-execution config overrides	68
5.3.3.46	LVV-T2458 - Verify serialization of pre-flight results	69
5.3.3.47	LVV-T2451 - Verify ability to append to an existing repository .	70
5.3.3.48	LVV-T2453 - Verify creation of DatasetRef upon butler.put . . .	71
5.3.3.49	LVV-T2449 - Verify middleware writer configurability	71
5.3.3.50	LVV-T2452 - Verify specification of output locations	72

5.3.3.51	LVV-T2450 - Verify writing dataset to multiple repositories . . .	73
5.3.3.52	LVV-T2447 - Verify DataRepository layering: Data Release and Science Platform	74
5.3.3.53	LVV-T2446 - Verify registries of collections	75
5.3.3.54	LVV-T2444 - Verify dataset garbage collection	76
5.3.3.55	LVV-T2442 - Verify dataset deletion	78
5.3.3.56	LVV-T2443 - Verify repository removal	79
5.3.3.57	LVV-T2441 - Verify repository version migration	80
5.3.3.58	LVV-T2440 - Verify versioning of DataRepositories	82
5.3.3.59	LVV-T2439 - Verify relocatability of DataRepositories	83
A Documentation		85
B Acronyms used in this document		85

LDM-GEN3: Gen 3 Butler Acceptance Testing Test Plan and Report

1 Introduction

1.1 Objectives

The goal of this test is to demonstrate that the Gen3 Butler software project has sufficiently matured that subsequent DM development should begin focusing on adoption of Gen3 Butler software repositories throughout the DM software project (i.e. that deprecation of Gen2 Butler usage within the project can begin).

1.2 System Overview

The Gen3 refactoring of the Butler is central to evolution of the overall DM software design and has repercussions throughout the rest of the DM project. This test plan is designed to verify that minimal requirements have been met and the DM project can now begin the process of integrating the Gen3 Butler within the pipelines and analysis tools. Those minimal requirements are that:

1. possible to ingest raw dataset types central to the Rubin operations and the ongoing development of the data management systems..
2. cp_pipe equivalent under Gen3 is available
3. developers can run a pipeline with a single-node using pipetask
4. processing supporting development is possible in a reasonable time (e.g. a 3-tract RC2 test run can be accomplished within a reasonable time)
5. Calibration Product Pipelines (CPP) can be run to support above investigations
6. Batch Processing System (BPS) is available to support testing at larger scales

In addition, at the time these tests occur the Gen3 Butler schema be considered stable enough that changes no longer occur on a weekly basis (i.e forced re-ingestion/migration of existing repositories are no longer a weekly occurrence). Changes requiring wholesale reingestion/migration may still be required but will occur in a regimented manner and the choice to allow

schema changes without an accompanying means to migrate old repositories would become a change-control board (CCB) level issue.

Applicable Documents:

LDM-592: Data Access Use Cases

LDM-556: Data Management Middleware Requirements

LDM-639: Data Management Acceptance Test Specification

1.3 Document Overview

This document was generated from Jira, obtaining the relevant information from the LVV-P77 Jira Test Plan and related Test Cycles (LVV-C160 LVV-C162 LVV-C190).

Section 1 provides an overview of the test campaign, the system under test (Software Products), the applicable documentation, and explains how this document is organized. Section 2 provides additional information about the test plan, like for example the configuration used for this test or related documentation. Section 3 describes the necessary roles and lists the individuals assigned to them.

Section 4 provides a summary of the test results, including an overview in Table 2, an overall assessment statement and suggestions for possible improvements. Section 5 provides detailed results for each step in each test case.

The current status of test plan LVV-P77 in Jira is **Draft** .

1.4 References

[1] **[DMTN-140]**, Comoretto, G., 2021, *Documentation Automation for the Verification and Validation of Rubin Observatory Software*, DMTN-140, URL <https://dmtn-140.lsst.io/>,

Vera C. Rubin Observatory Data Management Technical Note

- [2] **[DMTN-178]**, Comoretto, G., 2021, *Docsteady Usecases for Rubin Observatory Constructions*, DMTN-178, URL <https://dmtn-178.lsst.io/>,
Vera C. Rubin Observatory Data Management Technical Note
- [3] **[LDM-556]**, Dubois-Felsmann, G., Jenness, T., Bosch, J., et al., 2018, *Data Management Middleware Requirements*, LDM-556, URL <https://ldm-556.lsst.io/>,
Vera C. Rubin Observatory Data Management Controlled Document
- [4] **[LDM-639]**, Guy, L., Wood-Vasey, W., Bellm, E., et al., 2020, *LSST Data Management Acceptance Test Specification*, LDM-639, URL <https://ldm-639.lsst.io/>,
Vera C. Rubin Observatory Data Management Controlled Document
- [5] **[LDM-592]**, Jenness, T., Bosch, J., Gower, M., et al., 2018, *Data Access Use Cases*, LDM-592, URL <https://ldm-592.lsst.io/>,
Vera C. Rubin Observatory Data Management Controlled Document
- [6] **[LSE-160]**, Selvy, B., 2013, *Verification and Validation Process*, LSE-160, URL <https://lsst.org/lse-160>

2 Test Plan Details

2.1 Data Collection

Observing is not required for this test campaign.

2.2 Verification Environment

These tests assume a stable weekly stack which supports Gen3 running of the above, that services that automatically ingest new data can support on-going ingestion to Gen3 repositories (i.e. DBB shared spaces and OODS support serving data through Gen3), and that batch processing services can support pipeline execution of Gen3 products.

2.3 Related Documentation

No additional documentation provided.

2.4 PMCS Activity

Primavera milestones related to the test campaign:

- LDM-GEN3

3 Personnel

The personnel involved in the test campaign is shown in the following table.

T. Plan LVV-P77 owner: Robert Gruendl			
T. Cycle LVV-C160 owner: Robert Gruendl			
Test Cases	Assigned to	Executed by	Additional Test Personnel
LVV-T2264	Leanne Guy		
LVV-T1984	Leanne Guy		
LVV-T1982	Leanne Guy		
LVV-T1987	Leanne Guy		
LVV-T1983	Leanne Guy		
T. Cycle LVV-C162 owner: Undefined			
Test Cases	Assigned to	Executed by	Additional Test Personnel
LVV-T1985	Leanne Guy		
T. Cycle LVV-C190 owner: Jeffrey Carlin			
Test Cases	Assigned to	Executed by	Additional Test Personnel
LVV-T2503	Leanne Guy		
LVV-T2502	Leanne Guy		
LVV-T2501	Leanne Guy		
LVV-T2499	Leanne Guy		
LVV-T2498	Leanne Guy		
LVV-T2497	Leanne Guy		
LVV-T2496	Leanne Guy		
LVV-T2495	Leanne Guy		
LVV-T2494	Leanne Guy		
LVV-T2493	Leanne Guy		
LVV-T2492	Leanne Guy		
LVV-T2491	Leanne Guy		
LVV-T2488	Leanne Guy		
LVV-T2487	Leanne Guy		
LVV-T2486	Leanne Guy		
LVV-T2485	Leanne Guy		
LVV-T2484	Leanne Guy		
LVV-T2483	Leanne Guy		
LVV-T2482	Leanne Guy		
LVV-T2481	Leanne Guy		
LVV-T2480	Leanne Guy		
LVV-T2479	Leanne Guy		

LWV-T2478	Leanne Guy
LWV-T2477	Leanne Guy
LWV-T2476	Leanne Guy
LWV-T2474	Leanne Guy
LWV-T2475	Leanne Guy
LWV-T2473	Leanne Guy
LWV-T2472	Leanne Guy
LWV-T2471	Leanne Guy
LWV-T2470	Leanne Guy
LWV-T2469	Leanne Guy
LWV-T2468	Leanne Guy
LWV-T2466	Jeffrey Carlin
LWV-T2467	Leanne Guy
LWV-T2464	Leanne Guy
LWV-T2465	Jeffrey Carlin
LWV-T2461	Leanne Guy
LWV-T2463	Jeffrey Carlin
LWV-T2462	Jeffrey Carlin
LWV-T2460	Jeffrey Carlin
LWV-T2457	Jeffrey Carlin
LWV-T2456	Jeffrey Carlin
LWV-T2455	Jeffrey Carlin
LWV-T2454	Jeffrey Carlin
LWV-T2458	Jeffrey Carlin
LWV-T2451	Jeffrey Carlin
LWV-T2453	Jeffrey Carlin
LWV-T2449	Jeffrey Carlin
LWV-T2452	Jeffrey Carlin
LWV-T2450	Jeffrey Carlin
LWV-T2447	Leanne Guy
LWV-T2446	Leanne Guy
LWV-T2444	Leanne Guy
LWV-T2442	Leanne Guy
LWV-T2443	Leanne Guy
LWV-T2441	Leanne Guy
LWV-T2440	Leanne Guy
LWV-T2439	Leanne Guy

4 Test Campaign Overview

4.1 Summary

T. Plan LVV-P77:		LDM-GEN3: Gen 3 Butler Acceptance Testing		Draft
T. Cycle LVV-C160:		LDM-503-GEN3: Gen 3 Butler Acceptance Testing		Not Executed
Test Cases	Ver.	Status	Comment	Issues
LVV-T2264	1	Not Executed		
LVV-T1984	1	Not Executed		
LVV-T1982	1	Not Executed		
LVV-T1987	1	Not Executed		
LVV-T1983	1	Not Executed		
T. Cycle LVV-C162:		LDM-503-GEN3: Gen 3 Ingest raw dataset		Not Executed
Test Cases	Ver.	Status	Comment	Issues
LVV-T1985	1	Not Executed		
T. Cycle LVV-C190:		LDM-556: Middleware Acceptance Testing		Not Executed
Test Cases	Ver.	Status	Comment	Issues
LVV-T2503	1	Not Executed		
LVV-T2502	1	Not Executed		
LVV-T2501	1	Not Executed		
LVV-T2499	1	Not Executed		
LVV-T2498	1	Not Executed		
LVV-T2497	1	Not Executed		
LVV-T2496	1	Not Executed		
LVV-T2495	1	Not Executed		
LVV-T2494	1	Not Executed		
LVV-T2493	1	Not Executed		
LVV-T2492	1	Not Executed		
LVV-T2491	1	Not Executed		
LVV-T2488	1	Not Executed		
LVV-T2487	1	Not Executed		
LVV-T2486	1	Not Executed		
LVV-T2485	1	Not Executed		
LVV-T2484	1	Not Executed		
LVV-T2483	1	Not Executed		
LVV-T2482	1	Not Executed		
LVV-T2481	1	Not Executed		
LVV-T2480	1	Not Executed		

LWV-T2479	1	Not Executed
LWV-T2478	1	Not Executed
LWV-T2477	1	Not Executed
LWV-T2476	1	Not Executed
LWV-T2474	1	Not Executed
LWV-T2475	1	Not Executed
LWV-T2473	1	Not Executed
LWV-T2472	1	Not Executed
LWV-T2471	1	Not Executed
LWV-T2470	1	Not Executed
LWV-T2469	1	Not Executed
LWV-T2468	1	Not Executed
LWV-T2466	1	Not Executed
LWV-T2467	1	Not Executed
LWV-T2464	1	Not Executed
LWV-T2465	1	Not Executed
LWV-T2461	1	Not Executed
LWV-T2463	1	Not Executed
LWV-T2462	1	Not Executed
LWV-T2460	1	Not Executed
LWV-T2457	1	Not Executed
LWV-T2456	1	Not Executed
LWV-T2455	1	Not Executed
LWV-T2454	1	Not Executed
LWV-T2458	1	Not Executed
LWV-T2451	1	Not Executed
LWV-T2453	1	Not Executed
LWV-T2449	1	Not Executed
LWV-T2452	1	Not Executed
LWV-T2450	1	Not Executed
LWV-T2447	1	Not Executed
LWV-T2446	1	Not Executed
LWV-T2444	1	Not Executed
LWV-T2442	1	Not Executed
LWV-T2443	1	Not Executed
LWV-T2441	1	Not Executed
LWV-T2440	1	Not Executed
LWV-T2439	1	Not Executed

Table 2: Test Campaign Summary

4.2 Overall Assessment

Not yet available.

4.3 Recommended Improvements

Not yet available.

Draft

5 Detailed Test Results

5.1 Test Cycle LVV-C160

Open test cycle *LDM-503-GEN3: Gen 3 Butler Acceptance Testing* in Jira.

Test Cycle name: LDM-503-GEN3: Gen 3 Butler Acceptance Testing

Status: Not Executed

This test cycle is meant to demonstrate that the Gen3 butler and associated database and pipeline interfaces have matured to the point where they can replace the Gen2 butler. The test cases outlined here:

1. use a series of modest pipeline executions to show that the Gen3 software can support all future pipeline development,
2. those pipeline executions also show that a batch processing system (BPS) is available to enable that processing, and
3. demonstrate through inspection that documentation for developers exists
4. confirm that pipeline developers do not know of blockers if all future development assumes Gen3 Butler.

5.1.1 Software Version/Baseline

Not provided.

5.1.2 Configuration

Gen3 Butler repositories with test data are available within DBB spaces. Weekly DM stack has Gen3 and BPS elements present for tests.

5.1.3 Test Cases in LVV-C160 Test Cycle

5.1.3.1 LVV-T2264 - Butler Gen3 maturity sufficient to support future pipeline development.

Version **1**. Open *LW-T2264* test case in Jira.

This test is meant to verify that Butler Gen3 maturity is sufficient to provide comparable (or better) pipeline capabilities and results to those available under Butler Gen2.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	

Expected Result	

Actual Result	

5.1.3.2 LVV-T1984 - Demonstrate documentation/examples of Gen3 usage and cp_pipe equivalent.

Version **1**. Open *LW-T1984* test case in Jira.

Demonstrate the existence of fundamental documentation necessary to aid Gen2 users with

the transition to Gen3 use.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	
Identify document(s), web-pages, archived presentations, or example notebooks that provide documentation and/or examples of Gen3 functionality.	

Expected Result	
Document reference(s) or URL(s) for such documentation.	

Actual Result	

5.1.3.3 LWW-T1982 - Run a pipeline on a single node using pipetask.

Version 1. Open *LWW-T1982* test case in Jira.

To show that individual users have the ability to run either locally (w/ sqlite) or generally (w/ Postgres) using Gen3 Butler infrastructure.

Preconditions:

This test requires that Gen3 Butler infrastructure and underlying pipets have been integrated. It further requires (in spirit) that gen3 schema stability has been reached to facilitate compar-

ison of pipeline results with further stack development can be compared.

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	
Setup stack, identify inputs, pipetask execution of standard ci_hsc run.	

Test Data	
ci_hsc raw repository within a Gen3 Butler repo	

Expected Result	
Pipeline executes standard reduction without failure.	

Actual Result	

5.1.3.4 LVV-T1987 - Run Calibration Products Processing (CPP)

Version 1. Open *LW-T1987* test case in Jira.

Demonstrate that basic calibration processing from Gen2 era has been enabled within Gen3 environment. This test is not concerned with large scales but merely demonstrates that Gen3 capability to generate calibration products (i.e. they are no longer required to be generated in Gen2 and then migrated to Gen3).

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
<hr/>	
Description	
Identify an existing or instantiate a new Gen3 repository with raw bias, dark, and flat observations.	
<hr/>	
Test Data	
It is preferred these data be early observatory products (i.e. either AuxTel/LATISS or ComCam).	
<hr/>	
Expected Result	
A Gen3 repo with appropriate raw data products.	
<hr/>	
Actual Result	

Step 2	Step Execution Status: Not Executed
<hr/>	
Description	
Create master bias, dark and flat products from the raw products.	
<hr/>	
Expected Result	
A master bias, dark, and flat calibration product.	
<hr/>	
Actual Result	

5.1.3.5 LVV-T1983 - Mini RC2 processing capability

Version **1**. Open *LVV-T1983* test case in Jira.

Demonstrate that a typical 3-tract RC2 data processing is possible using the Gen3 system and the nascent Batch Production Service (BPS). This test is meant to demonstrate that Gen3 + BPS systems are capable of supporting future DM development by demonstrating that processing routinely used by developers for benchmarking/testing improvements can be performed in a reasonable time.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	
setup environment	
identify input data products	

Test Data	
RC2 raw repo	

Example Code	
# for example on lsstdev-* resources at NCSA	
<pre>export lsstsw_root=/software/lsstsw/stack export EUPS_TAG="w_2020_46"</pre>	

```
source /opt/rh/devtoolset-8/enable
source ${lsstsw_root}/loadLSST.bash
setup lsst_distrib -t ${EUPSTAG}
```

Expected Result

software environment ready for job submission

Actual Result

Step 2 Step Execution Status: **Not Executed**

Description

BPS pipeline submission

Example Code

```
pipetask qgraph -d "tract = 9615 and instrument='HSC' and skymap='hsc_rings_v1'" \
-b {gen3_repo}/{version}/butler.yaml \
-i HSC/calib,HSC/raw/all,HSC/masks,refcats,skymaps \
-p /home/madamow/gen2-to-gen3/bps/HSC-RC2.yaml \
-q /home/madamow/gen2-to-gen3/bps/submit/RC2/w_2020_42/DM-27244/20201102T10h22m03s/RC2_w_2020_42_DM-27244_20201102T10h22m03s.pickle
```

Expected Result

Pipeline execution is successful. An estimate of the compute resources used (# cores, memory, wall time) for each execution should be reported.

Actual Result

5.2 Test Cycle LVV-C162

Open test cycle *LDM-503-GEN3: Gen 3 Ingest raw dataset* in Jira.

Test Cycle name: LDM-503-GEN3: Gen 3 Ingest raw dataset

Status: Not Executed

In the context of the milestone LDM-503-GEN3, Gen 3 Butler readiness, this test cycle is defining the configuration and the dataset for running a generic **Raw Data Ingestion Into Gen3 Butler** test case. There are currently 5 data sources that require verification as they are the central products that will be produced by Rubin or are used as precursor sets in the development/verification of the data management software systems. The current raw data products that are deemed central to DM development and testing are those from AuxTel/LATISS, ComCam, and precursor data from HyperSuprimeCam (HSC). Note, further tests using LSSTCam (currently only preliminary BOT data from the SLAC test stand are available) or precursor sets from the Dark Energy Camera (DECam) could be added but since these types do not exactly fit the central model used for LSST they are not tied directly to requirements.

5.2.1 Software Version/Baseline

LSST DM Stack with Gen3 Butler.

5.2.2 Configuration

Three separate raw data types, those from: AuxTel/LATISS, ComCam, and HSC (e.g. a CI_HSC raw) should be ingested when this test is executed.

5.2.3 Test Cases in LVV-C162 Test Cycle

5.2.3.1 LVV-T1985 - Verify daf_butler raw data ingest

Version 1. Open *LW-T1985* test case in Jira.

Demonstrate that a raw data type can be successfully ingested into a Butler repository.

Preconditions:

In order to run this test, a Gen3 daf butler should be deployed and ready to use, with access to the filesystems where the raw data to ingest is stored.

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	
Identify data for ingestion HSC RC2 and make a copy at a location for the test. While a suggestion is provided in /project/shared/hsc/COSMOS/2014-03-27/ for a location where such data can be found, the actual data used can be left to the discretion of the person(s) executing the test with the added suggestion that relatively recent data are more likely to reflect the current observatory system state.	

Test Data	
/project/shared/hsc/COSMOS/2014-03-27/	

Expected Result	
One or more raw data sets are identified and made available.	

Actual Result	

Step 2	Step Execution Status: Not Executed
Description	
Identify data for ingestion AuxTel/LATISS and make a copy at a location for the test. While a suggestion is provided in /project/shared/auxTel/_parent/raw/2021-03-23/ for a location where such data can be found, the actual data used can be left to the discretion of the person(s) executing the test with the added suggestion that relatively recent data are more likely to reflect the current observatory system state.	

Test Data

/project/shared/auxTel/_parent/raw/2021-03-23/

Expected Result

One or more raw data sets are identified and made available.

Actual Result

Step 3 Step Execution Status: **Not Executed**

Description

Identify data for ingestion ComCam and make a copy at a location for the test. While a suggestion is provided in /project/shared/comCam/_parent/raw/2021-05-14/2021051400003/ for a location where such data can be found, the actual data used can be left to the discretion of the person(s) executing the test with the added suggestion that relatively recent data are more likely to reflect the current observatory system state.

Test Data

/project/shared/comCam/_parent/raw/2021-05-14/2021051400003/

Expected Result

One or more raw data sets are identified and made available.

Actual Result

Step 4 Step Execution Status: **Not Executed**

Description

Verify that Butler repository is available for the HSC RC2 (Note this needs to be a test repository rather than the central repository as the raw data should not already be present in the repository for this test.)

Test Data

url 1

Example Code

create empty Gen3 repo (for ComCam data)

butler create repo

butler register-instrument repo lsst.obs.lsst.LsstComCam

Expected Result

Actual Result

Step 5 Step Execution Status: **Not Executed**

Description

Verify that Butler repository is available for the AuxTel/LATISS (Note this needs to be a test repository rather than the central repository as the raw data should not already be present in the repository for this test.)

Test Data

url 2

Example Code

create empty Gen3 repo (for ComCam data)

butler create repo

butler register-instrument repo lsst.obs.lsst.LsstComCam

Expected Result

Actual Result

Step 6 Step Execution Status: **Not Executed**

Description

Verify that Butler repository is available for the ComCam (Note this needs to be a test repository rather than the central repository as the raw data should not already be present in the repository for this test.)

Test Data

url 3

Example Code

create empty Gen3 repo (for ComCam data)

```
butler create repo
butler register-instrument repo lsst.obs.lsst.LsstComCam
```

Expected Result

Actual Result

Step 7 Step Execution Status: **Not Executed**

Description

Ingest HSC RC2 raw data into repo

Test Data

url 1

Example Code

```
butler ingest-raws repo raw
```

Expected Result

Tool reports data ingest successful for HSC RC2 into url 1

Actual Result

Step 8 Step Execution Status: **Not Executed**

Description

Ingest AuxTel/LATISS raw data into repo

Test Data

url 2

Example Code

butler ingest-raws repo raw

Expected Result

Tool reports data ingest successful for AuxTel/LATISS into url 2

Actual Result

Step 9 Step Execution Status: **Not Executed**

Description

Ingest ComCam raw data into repo

Test Data

url 3

Example Code

butler ingest-raws repo raw

Expected Result

Tool reports data ingest successful for ComCam into url 3

Actual Result

Step 10 Step Execution Status: **Not Executed**

Description

Query repository to verify that ingestion of HSC RC2 occurred.

Test Data

url 1

Expected Result

HSC RC2 data are found by query.

Actual Result

Step 11 Step Execution Status: **Not Executed**

Description

Query repository to verify that ingestion of AuxTel/LATISS occurred.

Test Data

url 2

Expected Result

AuxTel/LATISS data are found by query.

Actual Result

Step 12 Step Execution Status: **Not Executed**

Description

Query repository to verify that ingestion of ComCam occurred.

Test Data

url 3

Expected Result

ComCam data are found by query.

Actual Result

5.3 Test Cycle LVV-C190

Open test cycle *LDM-556: Middleware Acceptance Testing* in Jira.

Test Cycle name: LDM-556: Middleware Acceptance Testing

Status: Not Executed

5.3.1 Software Version/Baseline

Not provided.

5.3.2 Configuration

Not provided.

5.3.3 Test Cases in LVV-C190 Test Cycle

5.3.3.1 LVV-T2503 - Verify Outouts from Test Processing Runs

Version **1**. Open *LW-T2503* test case in Jira.

Verify that the Data Output System interface is usable by algorithmic code being run for test/development purposes, on both development compute environments at the archive center and in personal environments.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
<hr/>	
Description	
Demonstrated by regular reprocessing runs at NCSA.	
<hr/>	
Expected Result	
<hr/>	
Actual Result	

5.3.3.2 LVV-T2502 - Verify Outputs from Science Platform

Version **1**. Open *LW-T2502* test case in Jira.

Verify that the Data Output System interface shall be usable by algorithmic code run in the Science Platform

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	
Demonstrated by any Science Platform notebook that uses the butler.	

Expected Result	

Actual Result	

5.3.3.3 LVV-T2501 - Verify Outputs from Data Release Production

Version 1. Open *LW-T2501* test case in Jira.

Verify that the Data Output System interface is usable by algorithmic code being run as part of Data Release Production.

Demonstrated by regular reprocessing runs at NCSA and DP0.2 production.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	

Expected Result	

Actual Result	

5.3.3.4 LVV-T2499 - Verify Consistent Output Interface

Version 1. Open *LW-T2499* test case in Jira.

Verify that the Data Output System provides a consistent interface for writing InMemory-Datasets to storage given a DatasetRef across different types of DataRepositories.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	

Expected Result	

Actual Result

5.3.3.5 LVV-T2498 - Verify Writing FITS tables

Version 1. Open *LVV-T2498* test case in Jira.

Verify that the Data Output System is able to write in-memory table objects as FITS files.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	

Expected Result

Actual Result

5.3.3.6 LVV-T2497 - Verify Writing FITS images

Version **1**. Open *LW-T2497* test case in Jira.

Verify that the Data Output System is be able to write in-memory image objects as FITS files

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	

Expected Result	

Actual Result	

5.3.3.7 LVV-T2496 - Verify filename invariance

Version **1**. Open *LW-T2496* test case in Jira.

Verify that for all datasets stored with unique filenames (or paths) as part of a Data Release, the name of the file retrieved by an external user is also unique and has a predictable name that is not dependent on data access mechanism

This behavior is not guaranteed by code, but it is the way we have configured our filename templates.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	

Expected Result	

Actual Result	

5.3.3.8 LVV-T2495 - Verify Combining composite datasets for export

Version **1**. Open *LVV-T2495* test case in Jira.

Verify that a facility is available to combine file-based composite datasets into a single file in a Scientific Data Format

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	

Expected Result	

Actual Result	

5.3.3.9 LVV-T2494 - Verify Strong exception guarantee

Version 1. Open *LVV-T2494* test case in Jira.

Verify that a put operation on the Data Output System provides the strong exception guarantee. If a put operation fails the previous state shall be restored.

This is the usual behavior, and we regard it as a bug when it is violated, and we don't currently have any known bugs of this type.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	

Expected Result	

Actual Result	

5.3.3.10 LVV-T2493 - Verify No clobber

Version 1. Open *LVV-T2493* test case in Jira.

Verify that it is possible to configure the Data Output System such that it is an error to attempt to persist a dataset that is already present in the output repository

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
--------	--

Description

Expected Result

Actual Result

5.3.3.11 LVV-T2492 - Verify Blocked write operation

Version 1. Open *LW-T2492* test case in Jira.

Verify that a put operation on the Data Output System blocks until it has either worked or failed

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1 Step Execution Status: **Not Executed**

Description

Expected Result

Actual Result

5.3.3.12 LVV-T2491 - Verify Creation of new DatasetTypes

Version 1. Open *LW-T2491* test case in Jira.

Verify that the Data Output system allows a new DatasetType to be registered with a DataRepository, programmatically and at Supertask preflight-time, allowing Datasets of that DatasetType to be added to that DataRepository thereafter

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
--------	--

Description

Expected Result

Actual Result

5.3.3.13 LVV-T2488 - Verify access outputs from test processing runs

Version 1. Open *LVV-T2488* test case in Jira.

Verify that the Data Input System shall provide access to processing runs initiated for test/de-development purposes, from the same compute environment in which the processing was run

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1 Step Execution Status: **Not Executed**

Description

Instantiate a Butler at NCSA or SLAC targeting a test run in '/repo/main'.

Expected Result

Actual Result

Step 2 Step Execution Status: **Not Executed**

Description

Call Butler.get.

Expected Result

Actual Result

Step 3 Step Execution Status: **Not Executed**

Description

Verify that data is correctly retrieved

Expected Result

Actual Result

5.3.3.14 LVV-T2487 - Verify Accessing official Data Releases

Version 1. Open *LVV-T2487* test case in Jira.

Verify that the Data Input System interface shall provide access to official Data Releases from the LSST Science Platform.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1 Step Execution Status: **Not Executed**

Description

Instantiate a butler on RSP targeting DP0.x collections.

Expected Result

Actual Result

Step 2 Step Execution Status: **Not Executed**

Description
Call 'Butler.get'

Expected Result

Actual Result

Step 3 Step Execution Status: **Not Executed**

Description
Verify that data is retrieved

Expected Result

Actual Result

5.3.3.15 LVV-T2486 - Verify Consistent input interface

Version 1. Open *LW-T2486* test case in Jira.

Verify that the Data Input System provides a consistent interface for loading Datasets into memory given a DatasetRef across different types of DataRepositories

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	
Run a 'PipelineTask' against a local SQLite+POSIX repo	

Expected Result	

Actual Result	

Step 2	Step Execution Status: Not Executed
Description	
Run the same 'PipelineTask' against a PostgreSQL+POSIX repo.	

Expected Result	

Actual Result	

Step 3 Step Execution Status: **Not Executed**

Description

Run the same 'PipelineTask' against a PostgreSQL+S3 repo.

Expected Result

Actual Result

5.3.3.16 LVV-T2485 - Verify Local caching of remote resources

Version 1. Open *LW-T2485* test case in Jira.

Verify that it is possible to configure the Data Input System to cache a local version of a Dataset that has been retrieved from a remote DataRepository.

Note that this doesn't really look distinct from DMS-MWBT-REQ-0055 anymore; I think 0055 was perhaps supposed to be some kind of shared-filesystem proxy for something that lives on even slower storage, like tape.

The specs are similar enough that the same test can be used

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1 Step Execution Status: **Not Executed**

Description

Enable datastore caching in a Butler client in RSP (or any S3-backed repo).

Expected Result

Actual Result

Step 2 Step Execution Status: **Not Executed**

Description

Run butler.get twice, check (e.g. trace logs) that the second comes from

Expected Result

Actual Result

5.3.3.17 LVV-T2484 - Verify Local proxy

Version 1. Open *LVV-T2484* test case in Jira.

Verify that it is possible to configure the Data Input system to use a local proxy to share remote retrievals of common Datasets

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	
Enable datastore caching in a Butler client in RSP (or any S3-backed repo).	

Expected Result	

Actual Result	

Step 2	Step Execution Status: Not Executed
Description	
Run 'butler.get' twice, check (e.g. trace logs) that the second comes from	

Expected Result	

Actual Result	

5.3.3.18 LVV-T2483 - Verify Failure on missing input file

Version 1. Open *LVV-T2483* test case in Jira.

Verify that it is possible via configuration to require the Data Input System to fail if an expected file is not found at the specified location

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	
Manually create QG with execution butler.	

Expected Result	

Actual Result	

Step 2	Step Execution Status: Not Executed
Description	
Run 'butler.get' against execution butler for ref that does not exist.	

Expected Result	

Actual Result	

5.3.3.19 LVV-T2482 - Verify Enabling PipelineTasks to execute

Version 1. Open *LVV-T2482* test case in Jira.

Verify that is possible for the Data Input System to construct a InMemoryDataset from a set

of files stored locally on disk (without a remote database connection)

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	
Manually create QG with execution butler	

Expected Result	

Actual Result	

Step 2	Step Execution Status: Not Executed
Description	
Run 'butler.get' against execution butler.	

Expected Result	

Actual Result	

5.3.3.20 LVV-T2481 - Verify third party datasets

Version 1. Open *LW-T2481* test case in Jira.

Verify that it is possible for the Data Input System to read from catalogs provided by outside sources using the same interface used for reading first class LSST datasets via a different plugin.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	
Make an empty repo.	

Expected Result

Actual Result

Step 2	Step Execution Status: Not Executed
Description	
Ingest some external parquet or FITS catalog.	

Expected Result

Actual Result

Step 3 Step Execution Status: **Not Executed**

Description

Call butler put

Expected Result

Actual Result

5.3.3.21 LVV-T2480 - Verify Item from Composite Datasets

Version 1. Open *LVV-T2480* test case in Jira.

Verify that it is possible to load into memory an item from a Composite Dataset without loading the full Dataset.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1 Step Execution Status: **Not Executed**

Description

Use a butler to read a component (e.g. WCS), against any repo.

Expected Result

Actual Result

5.3.3.22 LVV-T2479 - Verify Parameterized subset of a Dataset

Version **1**. Open *LW-T2479* test case in Jira.

Verify that It is possible to load into memory a parameterized subset of a Dataset without loading the full Dataset.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1 Step Execution Status: **Not Executed**

Description

Use a butler to read a subimage via get parameters against any repo.

Expected Result

Actual Result

5.3.3.23 LVV-T2478 - Verify I/O using cloud storage

Version 1. Open *LVV-T2478* test case in Jira.

Verify that the Data Input/Output System shall be able to utilize cloud-based storage engines.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1 Step Execution Status: **Not Executed**

Description

Make an empty repo with an S3 datastore,

Expected Result

Actual Result

Step 2 Step Execution Status: **Not Executed**

Description

Run 'butler get/put'

Expected Result

Actual Result

5.3.3.24 LVV-T2477 - Verify I/O using distributed file system

Version 1. Open *LW-T2477* test case in Jira.

Verify that the Data Input/Output System shall be able to read/write from/to distributed file systems.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1 Step Execution Status: **Not Executed**

Description

Make an empty repo with a POSIX datastore

Expected Result

Actual Result

Step 2 Step Execution Status: **Not Executed**

Description
do butler get/put

Expected Result

Actual Result

5.3.3.25 LVV-T2476 - Verify Format Plugability

Version 1. Open *LVV-T2476* test case in Jira.

Verify that it is possible to control the method used to read and write a particular DatasetType using a text configuration file such that the Python object and the form of the persisted dataset can be configured externally.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1 Step Execution Status: **Not Executed**

Description

Make an empty repo with the default configuration.

Expected Result

Actual Result

Step 2 Step Execution Status: **Not Executed**

Description

Make an empty repo with configuration that overrides a formatter.

Expected Result

Actual Result

Step 3 Step Execution Status: **Not Executed**

Description

Make an empty repo with configuration that changes a StorageClass's Python

Expected Result

Actual Result

Step 4 Step Execution Status: **Not Executed**

Description

Put and get the same datasets to all repos.

Expected Result

Actual Result

5.3.3.26 LVV-T2474 - Verify Data Discovery for Data Release Production

Version 1. Open *LW-T2474* test case in Jira.

Verify that the Data Discovery System interface is usable when initiating processing for Data Release Production.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
--------	--

Description

Run QG generation for the DRP pipeline against any major repo (e.g. '/repo/main'). Same as for LVV-T2473

Expected Result

Actual Result

5.3.3.27 LVV-T2475 - Verify Data discovery for test processing runs

Version **1**. Open *LVV-T2475* test case in Jira.

Verify that the Data Discovery System interface is usable when initiating processing runs initiated for test/development purposes (on LSST or personal hardware),

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	
Run QG generation against any major repo (e.g. '/repo/main'). Same script as for	

Expected Result	

Actual Result	

5.3.3.28 LVV-T2473 - Verify Consistent discovery interface

Version **1**. Open *LVV-T2473* test case in Jira.

Verify that the Data Discovery System provides a consistent interface for obtaining a graph

that represents the DataUnits and Datasets in a DataRepository that match user specified criteria.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	
Run QG generation against any major repo (e.g. '/repo/main').	

Expected Result	

Actual Result	

5.3.3.29 LVV-T2472 - Verify Introspection for DatasetExpressions

Version **1**. Open *LW-T2472* test case in Jira.

Verify that the Data Discovery System allows for a DatasetExpression to be constructed interactively using introspection on the DataRepository schema

Note that the requirement talks about high-level interactive tooling, but description makes it clear that middleware is only responsible for exposing the introspection necessary to allow

that tooling to be written, and we do.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	
Print dimension metadata schema by walking through DimensionUniverse.	

Expected Result	

Actual Result	

5.3.3.30 LVV-T2471 - Verify Filter by non-DatasetRef Database Entries

Version 1. Open *LW-T2471* test case in Jira.

Verify that the Data Discovery System is able to filter search results based upon specified filters that need non-DatasetRef database entries

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
<hr/>	
Description	
Run 'butler query-datasets' against any major repo (e.g. '/repo/main'), with a WHERE expression involving some dimension metadata fields.	
<hr/>	
Expected Result	
<hr/>	
Actual Result	

5.3.3.31 LVV-T2470 - Verify Dataset overrides

Version 1. Open *LW-T2470* test case in Jira.

Verify that it is possible for an operator to configure the Data Discovery System to override certain Datasets with others before retrieval.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1 Step Execution Status: **Not Executed**

Description

Run 'butler query-datasets' against any major repo (e.g. '/repo/main'), with multiple input collections that contain the same unresolved DatasetRefs, with findFirst=True.

Expected Result

Actual Result

5.3.3.32 LVV-T2469 - Verify Multiple parallel input Collections

Version 1. Open *LVV-T2469* test case in Jira.

Verify that the Data Discovery System is able to locate Datasets from multiple input Collections in order to retrieve the same logical Dataset from them all.

This is to allow for comparison of the same data reduced with multiple different stacks.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1 Step Execution Status: **Not Executed**

Description

Run 'butler query-datasets' against any major repo (e.g. '/repo/main'), with multiple input collections that contain the same unresolved DatasetRefs, with findFirst=False

Expected Result

Actual Result

5.3.3.33 LVV-T2468 - Verify Multiple chained input Collections

Version 1. Open *LVV-T2468* test case in Jira.

Verify that the Data Discovery System is able treat multiple input Collections as a single coherent logical repository

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1 Step Execution Status: **Not Executed**

Description

Run 'butler query-datasets' against any major repo (e.g. '/repo/main') with multiple input collections that contain the same unresolved DatasetRefs, with findFirst=False

Expected Result

Actual Result

5.3.3.34 LVV-T2466 - Verify enable complete pipeline specification

Version 1. Open *LVV-T2466* test case in Jira.

Verify that the design provides an interface for delivering a complete algorithmic work specification (a “Pipeline specification”) from Science Pipelines to an execution system, the “supervisory framework”, a notable instance of which is the LSST production system.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
--------	--

Description

This is a fundamental part of the design of PipelineTask.

Expected Result

Actual Result

5.3.3.35 LVV-T2467 - Verify DataUnit lookup: processing driven

Version **1**. Open *LVV-T2467* test case in Jira.

Verify that all Data Discovery Systems make it possible to discover the DataUnits for all Datasets that could potentially be used to produce a given DatasetType with known DataUnits.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	
Run 'butler query-datasets' against any major repo (e.g. '/repo/main').	

Expected Result	

Actual Result	

5.3.3.36 LVV-T2464 - Verify multiple simultaneous sky definitions

Version **1**. Open *LVV-T2464* test case in Jira.

Verify that a collection is able to hold Datasets corresponding to different sky tilings simulta-

neously

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	
Make empty repo.	

Expected Result	

Actual Result	

Step 2	Step Execution Status: Not Executed
Description	
Run 'butler register-skymap'.	

Expected Result	

Actual Result	

Step 3 Step Execution Status: **Not Executed**

Description

Run 'butler register-skymap' a second time

Expected Result

Actual Result

Step 4 Step Execution Status: **Not Executed**

Description

Verify that mappings to both tile definitions are valid

Expected Result

Actual Result

5.3.3.37 LVV-T2465 - Verify pipeline execution in multiple contexts

Version 1. Open *LVV-T2465* test case in Jira.

Verify that the design allows a given Pipeline specification to be used in both development and production contexts.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	
This is a fundamental part of the design of PipelineTask.	

Expected Result	

Actual Result	

5.3.3.38 LVV-T2461 - Verify Collection Layering: Science Platform

Version 1. Open *LW-T2461* test case in Jira.

Verify that collections created in the Science Platform are usable as inputs for processing initiated in the Science Platform

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	

Run part of DRP pipeline in RSP.

Expected Result

Actual Result

Step 2 Step Execution Status: **Not Executed**

Description

Run a later part of DRP pipeline in RSP

Expected Result

Actual Result

5.3.3.39 LVV-T2463 - Verify enabling of different execution environments

Version 1. Open *LW-T2463* test case in Jira.

Verify that the supervisory framework supports the creation of multiple specializations for different execution environments.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	
Satisfied by BPS plugin system; we have plugins for many workflow systems already.	

Expected Result	

Actual Result	

5.3.3.40 LVV-T2462 - Verify QuantumGraph algorithm

Version 1. Open *LVV-T2462* test case in Jira.

Verify QuantumGraph algorithm common to all execution environments. Verify that the supervisory framework provides a common implementation of the logic required for interpretation of the Pipeline steps and their data groupings (and thus the possible parallelization); i.e., that the QuantumGraph generation algorithm can be common to all execution environments.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	

Requirement is an attempt to require a QuantumGraph generation algorithm that is common to all execution environments, without actually saying QuantumGraph generation.

Expected Result

Actual Result

5.3.3.41 LVV-T2460 - Verify generating a DAG

Version 1. Open *LVV-T2460* test case in Jira.

Verify that the supervisory framework supports the “Pre-flight” phase of execution of a Pipeline on a specified set of inputs and/or desired outputs, resulting in a Directed Acyclic Graph (DAG) for the processing, with the nodes in the DAG being the units of work to be executed.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
--------	--

Description

Satisfied by existence of QuantumGraph generation code.

Expected Result

Actual Result

5.3.3.42 LVV-T2457 - Verify butler instantiation

Version 1. Open *LVV-T2457* test case in Jira.

Verify that the supervisory framework creates and supplies the Butler required to support the I/O to be performed in the "Run" phase, for each unit of work.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
--------	--

Description

Expected Result

Actual Result

5.3.3.43 LVV-T2456 - Verify execution logging

Version **1**. Open *LW-T2456* test case in Jira.

Verify that standard logging is enabled for the pre-flight and run processes of pipelines.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	

Expected Result	

Actual Result	

5.3.3.44 LVV-T2455 - Verify pipeline interface available as Python API

Version **1**. Open *LW-T2455* test case in Jira.

Verify that the Pipeline specification interface is available as a Python API.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	

Expected Result	

Actual Result	

5.3.3.45 LVV-T2454 - Verify pre-execution config overrides

Version 1. Open *LVV-T2454* test case in Jira.

Verify that the middleware enables programmatic overrides to the configurations specified for a Pipeline, and that the overrides can be captured for purposes of provenance recording.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	

Expected Result	

Actual Result	

5.3.3.46 LVV-T2458 - Verify serialization of pre-flight results

Version 1. Open *LW-T2458* test case in Jira.

Verify that the supervisory framework provides a serialization form for the results of the “Pre-flight” phase, so that they can be computed in one process and executed under the control of one or more others.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	
Satisfied by QuantumGraph being serializable.	

Expected Result	

Actual Result

5.3.3.47 LVV-T2451 - Verify ability to append to an existing repository

Version 1. Open *LVV-T2451* test case in Jira.

Verify that it is possible to add Datasets to a pre-existing Collection via additional processing.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	
In regular use in production and included in unit tests. - WHICH UNIT TESTS?	

Expected Result

Unit test passes

Actual Result

5.3.3.48 LVV-T2453 - Verify creation of DatasetRef upon butler.put

Version 1. Open *LVV-T2453* test case in Jira.

Verify that upon writing a dataset, a DatasetRef is created to enable getting the dataset in the future.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	
Execute ButlerPutGetTests in https://github.com/lsst/daf_butler/blob/main/tests/test_butler.py	

Expected Result	
Unit test passes	

Actual Result	

5.3.3.49 LVV-T2449 - Verify middleware writer configurability

Version 1. Open *LVV-T2449* test case in Jira.

Verify that the data output system supports configuration of individual writer behavior.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	
Execute the unit test at https://github.com/lsst/daf_butler/blob/main/tests/test_config.py , which tests the writer configuration.	

Expected Result	
Unit test passes.	

Actual Result	

5.3.3.50 LVV-T2452 - Verify specification of output locations

Version **1**. Open *LW-T2452* test case in Jira.

Verify that the middleware enables configuration of the output location for a POSIX file system.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	
Execute the PosixDatastoreButlerTestCase in https://github.com/lsst/daf_butler/blob/main/tests/test_butler.py	

Expected Result	
Unit test passes	

Actual Result	

5.3.3.51 LVV-T2450 - Verify writing dataset to multiple repositories

Version 1. Open *LVV-T2450* test case in Jira.

Verify that the middleware enables writing of a single dataset to multiple repositories, with a different output format used for each repository.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
--------	--

Description

Not regularly exercised in production, but functionality is tested in daf_butler unit tests. — WHICH UNIT TESTS?

Expected Result

Unit test passes

Actual Result

5.3.3.52 LVV-T2447 - Verify DataRepository layering: Data Release and Science Platform

Version 1. Open *LVV-T2447* test case in Jira.

Verify that a Data Release is usable as the inputs for processing initiated in the Science Platform.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1 Step Execution Status: **Not Executed**

Description

Reuse test script for DMS-MWBT-REQ-0012

Expected Result

Actual Result

Step 2 Step Execution Status: **Not Executed**

Description

Run a DRP pipeline subset in RSP, using DP0.x collections as inputs.

Expected Result

Actual Result

5.3.3.53 LVV-T2446 - Verify registries of collections

Version 1. Open *LVV-T2446* test case in Jira.

Verify that there is a mechanism for registering Collections as they are created

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1 Step Execution Status: **Not Executed**

Description

Expected Result

Actual Result

5.3.3.54 LVV-T2444 - Verify dataset garbage collection

Version 1. Open *LVV-T2444* test case in Jira.

Verify that when a DataRepository is removed, the Datasets it references are removed if and only if they are not also referenced by one or more additional DataRepositories that have been explicitly identified.

Note that the requirement text assumed a slightly different collections model from what we have. Instead of “reference counting” datasets, we have RUN collections that own datasets and TAGGED collections that don’t, but we still guard against improper deletions as the requirement demands.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
--------	--

Description

Make example repo, one each, POSIX and S3

Expected Result

Actual Result

Step 2 Step Execution Status: **Not Executed**

Description

Create TAGGED collection and add some datasets to it.

Expected Result

Actual Result

Step 3 Step Execution Status: **Not Executed**

Description

Try to delete the RUN collection - shouldn't be possible because of references in TAGGED collection.

Expected Result

Actual Result

Step 4 Step Execution Status: **Not Executed**

Description

Try to delete the TAGGED collection - should work, without deleting the datasets.

Expected Result

Actual Result

5.3.3.55 LVV-T2442 - Verify dataset deletion

Version 1. Open *LVV-T2442* test case in Jira.

Verify that a Dataset is deletable from a DataRepository by an authorized perso

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1	Step Execution Status: Not Executed
Description	
Make example repo, one each, POSIX and S3	

Expected Result

Actual Result

Step 2	Step Execution Status: Not Executed
Description	
Run 'butler prune-datasets'	

Expected Result

Actual Result

Step 3 Step Execution Status: **Not Executed**

Description

Verify that the datasets are deleted

Expected Result

Actual Result

5.3.3.56 LVV-T2443 - Verify repository removal

Version **1**. Open *LVV-T2443* test case in Jira.

Verify that an authorized user can remove a DataRepository from any storage environment. Verification on **all** environments is not possible We will verify POSIX and S3 environments, which we believe is in the spirit of the requirement and covers our core needs.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1 Step Execution Status: **Not Executed**

Description

Make example repo, one each, POSIX and S3

Expected Result

Actual Result

Step 2 Step Execution Status: **Not Executed**

Description

Run 'butler prune-collection'

Expected Result

Actual Result

Step 3 Step Execution Status: **Not Executed**

Description

Verify that the repository is removed

Expected Result

Actual Result

5.3.3.57 LVV-T2441 - Verify repository version migration

Version 1. Open *LW-T2441* test case in Jira.

Verify that the Data Input/Output system can perform persistent migrations of a DataRepository to bring the Data Model of that DataRepository up to parity with the Data Model expected by the current Data Input/Output System interfaces.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1 Step Execution Status: **Not Executed**

Description

Make example repo with one schema configuration (e.g. autoincrement dataset IDs).

Expected Result

Actual Result

Step 2 Step Execution Status: **Not Executed**

Description

Run existing migration script that upgrades the repository to the new one.

Expected Result

Actual Result

Step 3 Step Execution Status: **Not Executed**

Description

Verify that the features are migrated correctly

Expected Result

Actual Result

5.3.3.58 LVV-T2440 - Verify versioning of DataRepositories

Version 1. Open *LW-T2440* test case in Jira.

Verify that the Data Input/Output system can describe the version of a DataRepository

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1 Step Execution Status: **Not Executed**

Description

Print attributes table via Butler APIs, against literally any repo.

Expected Result

Actual Result

5.3.3.59 LVV-T2439 - Verify relocatability of DataRepositories

Version **1**. Open *LW-T2439* test case in Jira.

Verify that DataRepositories can be relocated between various storage contexts.

Preconditions:

Execution status: **Not Executed**

Final comment:

Detailed steps results:

Step 1 Step Execution Status: **Not Executed**

Description

Make empty repo with local SQLite+POSIX storage in RSP.

Expected Result

Actual Result

Step 2 Step Execution Status: **Not Executed**

Description

Use 'butler transfer-from' to transfer datasets from DP0.x repos to local repo, with transfer=copy.

Expected Result

Actual Result

Step 3 Step Execution Status: **Not Executed**

Description

Check that the relocated dataset is the same as the original

Expected Result

Actual Result

A Documentation

The verification process is defined in LSE-160. The use of Docsteady to format Jira information in various test and planing documents is described in DMTN-140 and practical commands are given in DMTN-178.

B Acronyms used in this document

Acronym	Description
API	Application Programming Interface
BOT	Bench for Optical Testing
BPS	Batch Production Service
CCB	Change Control Board
CPP	Calibration Production Processing
ComCam	The commissioning camera is a single-raft, 9-CCD camera that will be installed in LSST during commissioning, before the final camera is ready.
DBB	Data Backbone
DECam	Dark Energy Camera
DM	Data Management
DMS	Data Management Subsystem
DMTN	DM Technical Note
DPO	Data Preview 0
DRP	Data Release Production
FITS	Flexible Image Transport System
HSC	Hyper Suprime-Cam
LATISS	LSST Atmospheric Transmission Imager and Slitless Spectrograph
LDM	LSST Data Management (Document Handle)
LSE	LSST Systems Engineering (Document Handle)
LSST	Legacy Survey of Space and Time (formerly Large Synoptic Survey Telescope)
LVV	LSST Verification and Validation
NCSA	National Center for Supercomputing Applications
OODS	Observatory Operations Data Service
PMCS	Project Management Controls System

POSIX	Portable Operating System Interface
RSP	Rubin Science Platform
S3	(Amazon) Simple Storage Service
SLAC	SLAC National Accelerator Laboratory
URL	Universal Resource Locator
WCS	World Coordinate System
bps	bit(s) per second

Draft