

Gen3 Milestone FY2019-2

Running RC2 via Gen3 + Pegasus + Oracle

Not as a production operator*, but as a special developer who shares repo with other developers. Not delivering a polished, efficient service, but rather delivering the starting point for study and more work. Besides the presentation, trying to have something usable/maintainable for the following 3 months or so while next set of work is ongoing.

Deliverables

- Results of running Gen3 DRP pipeline on RC2 HSC data on lsst-dev
 - Using Oracle as backend for registry of shared butler repository
 - Pegasus for workflow
- Hsin-Fang can run Gen3 as part of the monthly HSC-RC2 reprocessing runs
 - Incorporate into procedures with lower expectations than Gen2 (failure rate, usefulness of outputs, operability, etc)
- Instructions for friendly-user developers using Gen3 Butler with Oracle
 - Including accessing weekly ci_hsc repo and RC outputs
 - Assuming running in Pegasus would still be too rough for most friendly-user developers
- Maintainable Registry code that does not require updating a copy of large portions of the code for each RDBMS for every code change

Why These Deliverables

1. Larger scale than ci_hsc
2. Oracle admins can start seeing data flow and more easily provide feedback
3. Increases visibility of Gen3 results
4. Gen2 RC running is one of the blockers to deprecating Gen2
5. Unblocks multi-registry work required for separation of production data from user data
6. Provides insight for the Batch Processing Service design doc

Delivery Date

DMLT: June 4-6, 2019

To do (Mid-level description to enable low level work and effort generation)

1. Design how to work in Oracle for demo and 3+ months following Oracle schemas
 - a. Distinct use cases will use distinct schemas (Registry developer, Pipetask developer, RC data user).
 - i. RC runs - Oracle admins owned and maintained.
 - ii. Weekly ci_hsc data loaded by Oracle admins into shared schema.
 - iii. More frequent schema updates - User owned and maintained.
 1. Need easy way to update schema with latest changes or start from scratch (Gen3) Schema evolution (schema migration scripts? (Gen3))
 2. Oracle Server setup (NCSA)
 - a. Recoverability - shared registries increase risk - Nightly exports of every schema will be scheduled with 2 week retention until data volumes warrant a different approach.
 - b. Availability - Standard maintenance windows and support during business hours.
 - c. Authentication - Oracle wallets (initially created by db admins).
 3. Install software on lsst-dev: Oracle client software + cx_Oracle (NCSA)
2. Increased testing:
 - a. More tests in PipelineTasks, ctrl_mpexec, etc (Jenkins + sqlite)
 - b. Running ci_hsc with Gen3 (sqlite) in Jenkins
 - c. Manual running of daf_butler tests against Oracle (pre-existing schema)
 - d. Stretch goals:
 - i. Jenkins running daf_butler tests against Oracle (full setup and teardown of schema)
 - ii. Running ci_hsc with Gen3 (oracle) in Jenkins
3. Have different output DataStores for different users (details TBD)
4. daf_butler refactoring work to decrease additional changes needed to function with multiple RDBMS products (Gen3).
5. Oracle specific Butler changes (NCSA + Gen3, blocked by refactoring work)
6. Need RC2 initial repo (Gen3)
 - a. (prefer) Ingest raw executable (+ script to make easier to start from scratch) (Gen3) Calib files may be ingest + script to set ranges
 - b. Or conversion from Gen2 HSC-RC2 reprocessing runs (like we do with ci_hsc) (ChrisW) Set initial WCS (only explicit update, not select best)
7. More Pipeline Tasks to convert to Gen3 (DRP)
 - a. SkyCorrection (needs to be broken up into smaller tasks)
 - b. JointCal (cannot be run on ci_hsc data set, needs more data)
8. Change template to have unique filenames for RC runs
 - a. Hopefully just saving the templates to a file. (NCSA)
 - b. Unknown if particular values in templates would require any Butler changes (Gen3)
9. Batch Processing Service - NCSA

- a. Assuming still using Andy's pipetask as the activator
 - b. Need execution config (in particular cpu/memory requirements)
 - c. Changes to allocNodes to set up HTCondor pool with partitionable slots
 - d. Helpful status/monitoring scripts TBD
 - e. Note the following are blocked by Gen3 development and are not part of this deliverable:
 - i. Must always start from beginning of submission (no retries or restarts)
 - ii. Must be shared repo model (no job scratch, no Pegasus file transfer)
10. RC2 dataset challenges:
- a. Single frame processing failures should not halt running
 - i. current proposed solution: config option to always write files
 - b. Missing warp file should not halt running
 - i. Ran into this with ci_hsc - config option exists to always write files
11. ci_hsc/RC2 output usable from NCSA LSP
- a. Oracle software accessible from NCSA LSP (NCSA + LSP/SQRE)
 - b. Not supporting Pegasus submissions from LSP for this milestone

* Why the note about not being Production?








- Missing separation of production data from user data (requires user write access to production schema)
- The outputs of a production pipeline should not be directly written to the production Data Backbone (or central database in general) to allow the Batch Production Service to:
 - Minimize database connections
 - Use various methods for retries and restarts
- Many missing Batch Production Service production features, some of which are blocked by not-yet-implemented Gen3 features.













Current Isst-dev Oracle Instructions









- For this milestone, no attempts will be made to make Oracle a part of the Isst_stack.
- Oracle instantclient and cx_oracle are currently installed on Isst-dev in */project/production/oracle*.
- Example Oracle environment settings are in */project/production/oracle/oracle_env-v1.sh*. Should not affect environment set up by LSST stack.
- Untar Oracle wallet tar given to you by admin in some directory in your home directory..
- The admin will also have given you a net service name for the wallet (e.g., *gen3_cred_yourlogin_1*). If provided the whole *tnsnames.ora* file, the net service name is the top/outer-most key.
- You will also need a *sqlnet.ora* and *tnsnames.ora* files (make sure the path in *sqlnet.ora* points to the wallet files).
- Set environment variable *TNS_ADMIN* to point to the directory where the *.ora files live.
- Test connection:
 - *sqlplus:*
`sqlplus /@<net service name>`
`select user from dual;`
`quit;`
 - Use test python program that prints who you connected to Oracle as if successful (Note: Must have python3 in your path. If you haven't already, `source /software/lsstsw/stack/loadLSST.bash`) :
`/project/production/oracle/test_conn.py <net service name>`

Timeline

Thurs Gen3 meet in gs	ci_hsc/RC2 Running	NCSA - BPS	NCSA - Oracle	Gen3	DRP
2 /2 1 /2 0 19	Completed ci_hsc gen2 run (sqlite to load into Oracle), ci_hsc gen3 run (sqlite3) in pegasus to provide feedback if things are no longer working.		Jim Bosch Oracle account		Must provide updated weekly Gen3 science configs prior to NCSA run
2 /2 8 /2 0 19	Completed ci_hsc gen2 run (sqlite to load into Oracle), ci_hsc gen3 run (sqlite3) in pegasus to provide feedback if things are no longer working.		Completed: Init Oracle accounts+wallets (Nate - 03/01), nightly DB backups (03/04), weekly ingest of ci_hsc, install Oracle client and cx_Oracle on Isst-dev (03/01)		

<div>3</div> <div>/7</div> <div>/2</div> <div>0</div> <div>19</div>	<div>Completed ci_hsc gen2 run (sqlite to load into Oracle)</div> <div>  DM-18176 - Jira project doesn't exist or you don't have permission to view it. </div> <div>  DM-18336 - Jira project doesn't exist or you don't have permission to view it. </div> <div>, ci_hsc gen3 run (sqlite3) in pegasus to provide feedback if things are no longer working.</div> <div>  DM-18177 - Jira project doesn't exist or you don't have permission to view it. </div>			<div>Completed Filename template checking script</div> <div>  </div> <div>DM-</div> <div>18181</div> <div>-</div> <div>Jira</div> <div>proje</div> <div>ct</div> <div>does</div> <div>n't</div> <div>exist</div> <div>or</div> <div>you</div> <div>don't</div> <div>have</div> <div>per</div> <div>miss</div> <div>ion</div> <div>to</div> <div>view</div> <div>it.</div>	
<div>3</div> <div>/1</div> <div>4</div> <div>/2</div> <div>0</div> <div>19</div>	<div>Completed ci_hsc gen2 run (sqlite to load into Oracle).</div> <div>  DM-18176 - Jira project doesn't exist or you don't have permission to view it. </div> <div>  DM-18336 - Jira project doesn't exist or you don't have permission to view it. </div> <div>ci_hsc gen3 run (sqlite3) in pegasus to provide feedback if things are no longer working.</div> <div>  DM-18178 - Jira project doesn't exist or you don't have permission to view it. </div>			<div>Decision about how to support multiple RDBMSs. Completed code changes for sqlite side. Code ready to start making Oracle changes.</div>	

3 /2 1 /2 0 19	<p>Completed ci_hsc gen2 run (sqlite to load into Oracle),</p> <div> DM-18176 - Jira project doesn't exist or you don't have permission to view it.</div> <div> DM-18336 - Jira project doesn't exist or you don't have permission to view it.</div> <p>ci_hsc gen3 run (sqlite3) in pegasus to provide feedback if things are no longer working.</p> <div> DM-18179 - Jira project doesn't exist or you don't have permission to view it.</div>	<p>Completed BPS v0.1 exec config, allocateNodes (partitionable slots),</p> <div> DM-18351 - Jira project doesn't exist or you don't have permission to view it.</div> <p>Completed unique filename templates ci_hsc (where only requires config file change)</p> <div> DM-18356 - Jira project doesn't exist or you don't have permission to view it.</div>	<p>Completed: Easy way to initialize dev butler schema in Oracle</p>	
3 /2 8 /2 0 19	<p>Completed ci_hsc gen2 run (sqlite to load into Oracle)</p> <div> DM-18176 - Jira project doesn't exist or you don't have permission to view it.</div> <div> DM-18336 - Jira project doesn't exist or you don't have permission to view it.</div> <p>, ci_hsc gen3 run (sqlite3) in pegasus to provide feedback if things are no longer working.</p> <div> DM-18180 - Jira project doesn't exist or you don't have permission to view it.</div>			
4 /4 /2 0 19	<p>Completed ci_hsc gen2 run (sqlite to load into Oracle)</p> <div> DM-18782 - Jira project doesn't exist or you don't have permission to view it.</div> <p>, ci_hsc gen3 run (sqlite3) in pegasus to provide feedback if things are no longer working.</p> <div> DM-18783 - Jira project doesn't exist or you don't have permission to view it.</div>		<p>Completed: Oracle Butler works (no efficiency checks, just doesn't abort). Selecting Oracle schema</p> <div> DM-18629 - Jira project doesn't exist or you don't have permission to view it.</div> <p>Table and view names case-insensitive on DB side.</p> <div> DM-17023 - Jira project doesn't exist or you don't have permission to view it.</div>	

4 /1 /2 0 19	<p>Completed ci_hsc gen2 run (sqlite to load into Oracle)</p> <div>  DM-18782 - Jira project doesn't exist or you don't have permission to view it. </div> <p>, ci_hsc gen3 run (sqlite) in pegasus to provide feedback if things are no longer working.</p> <div>  DM-18825 - Jira project doesn't exist or you don't have permission to view it. </div>	<p>Completed BPS v0.1 status/history scripts</p> <div>  DM-18780 - Jira project doesn't exist or you don't have permission to view it. </div>		<p>Completed: Scripts to initialize ci_hsc repo for a Gen3 run without latest weekly Gen2 outputs.</p>	
4 /1 /2 0 19	<p>Completed ci_hsc gen2 run (sqlite to load into Oracle)</p> <div>  DM-18782 - Jira project doesn't exist or you don't have permission to view it. </div> <p>, ci_hsc gen3 run (Oracle) in pegasus to provide feedback if things are no longer working.</p> <div>  DM-18830 - Jira project doesn't exist or you don't have permission to view it. </div>			<p>Completed: Mechanisms to create RC2 init repo</p>	
4 /2 /5 /2 0 19	<p>Completed ci_hsc gen2 run (sqlite to load into Oracle)</p> <div>  DM-18782 - Jira project doesn't exist or you don't have permission to view it. </div> <p>, ci_hsc gen3 run (Oracle) in pegasus to provide feedback if things are no longer working.</p> <div>  DM-18831 - Jira project doesn't exist or you don't have permission to view it. </div>		<p>Completed: RC2 init repo avail in Oracle</p> <div>  DM-18829 - Jira project doesn't exist or you don't have permission to view it. </div>		<p>Complete RC2 DRP pipeline includes always write output config options where needed.</p>
4 /2 /5 /2 0 19		<p>Freeze: features, API, schema</p>			

5 /2 /2 0 19	Start running RC2 and reporting problems				
5 /9 /2 0 19					
5 /1 6 /2 0 19					
5 /2 3 /2 0 19					
5 /3 0 /2 0 19			Completed: Can access Oracle Registry + GPFS DataStore from NCSA LSP		
6 /0 6 /2 0 19	Milestone completed. Presentation during DMLT meeting June 04-06. Includes instructions, any software installs, etc				