Large Synoptic Survey Telescope

Data Management Monthly Report

August 2017

High-level Summary

DM attended and supported the LSST2017 Project and Community Workshop. Many interesting sessions were held including commissioning simulations with the Commissioning Team. DM provided the stack in a JupyterLab environment for project and community members to try out.

We spent time preparing for the NSF/DOE Joint Status Review, to be held in September. This was, perhaps, easier for DM as we had all material ready from last month's DM Review. In the run up to the JSR, there was a Director's Review which was handled by the Deputy DMPM (Swinbank) and Jeff Kantor.

DM supported the LSST Community Science Center working group's face to face meeting at NOAO, where Melissa Graham gave a presentation on LSST data products, and contributed to writing the white paper report and William O'Mullane presented DM and the Science Platform.

We continued to focus on the development of the key "end-to-end" system for testing Alert Production pipelines. The "robust coaddition" algorithmic prototype has been developed into a task which is now being integrated into the DM Stack. We have completed initial experiments into measuring galaxy shears directly on coadded images. All remaining parts of Qserv, associated tooling and scripts, and the DAX web services were ported to Python 3, and all C++ wrappers were ported to pybind11. Result retrieval for asynchronous/disconnect queries was implemented in Qserv.

The build system has been improved allowing us to get better reports on test durations and status. FIrefly was deployed in a Docker container at NCSA. It is able to access the GPFS file system used for bulk data storage, making it possible to browse through re-processed HSC images at NCSA.

We made a proposal to Project Systems Engineering for a detailed schedule, with dependencies, for the early integration of Observatory subsystems with each other, leading up to operations rehearsals and the deployment of the auxiliary telescope spectrograph. The proposal was discussed, refined, and adopted in sessions at the PCW, and will be incorporated into the PMCS.

The Data Acquisition (DAQ) test stand deployment activity was completed with security vetting and follow-up in NCSA. Work on Disaster Recovery (DR) design and implementation

was completed, with verification of successful backup and restore of several partitions. The AURA Dense Wavelength Division Multiplexing (DWDM) equipment was received. Delays in REUNA DWDM equipment receipt from the vendor and in access to the Cerro Pachon caseta delayed the start of the equipment installation by one month. While this will not impact LSST commissioning, it will delay the First Optic Light Demonstration, as well as the installation of NCSA's Identity Management equipment, into November.

Risk Management

The DM Risk Register was reviewed in the monthly process. No new risks were added and no significant changes to existing risk exposure were made.

Milestone summary

Achieved Milestones

ID	Description	Due Date
	Science User Interface	
DLP-205	LSST web UI search external catalog	2017-11-30
DLP-207	Extra LSST query capabilities in web UI	2017-11-30
DLP-222	Provide time-domain data view	2017-11-30
	International Communications & Base Site	
DMTC-7400-2210	Fiber delivery by Telefonica	2017-02-28
	Science Quality & Reliability Engineering	
DLP-338	Release JUnit type testing framework and adapt stack	2016-05-31

Delayed Milestones

After the replan we still carry some DLP milestones which will be completed in FY17. It will take a few months to get the LCRs baselined and the milestones cleaned up.

ID	Description	Due Date
	System Management	
DLP-558	LSST Software Release 7.1 Complete	2017-08-31

	The release was delayed slightly but will be made in September.	
	DM System Architecture	
DLP-538, DLP-539, DLP-541, DLP-544	Assorted ICDs updated to Phase 3 *These are delayed milestones reported at the July 2017 DM Review. It is expected to be completed during the the F17 cycle. There is no impact on the critical path.*	2017-05-31
	Science Data Archive & Application Services	
DLP-802	AP/L1 Database Design *This is one of the delayed milestones reported at the July 2017 DM Review. It is expected to be completed during the the F17 cycle. There is no impact on the critical path.*	2016-11-30
DLP-472	Qserv Data Distribution *This is one of the delayed milestones reported at the July 2017 DM Review. It is expected to be completed during the the F17 cycle. There is no impact on the critical path.*	2017-05-31
	LSST Data Facility	
DLP-366	Mid-scale processing of eligible camera data *This is one of the delayed milestones reported at the July 2017 DM Review. It is expected to be completed during the the F17 cycle. There is no impact on the critical path.*	2016-11-30
DLP-412	Monitor development cluster *This is one of the delayed milestones reported at the July 2017 DM Review. It is expected to be completed during the the F17 cycle. There is no impact on the critical path.*	2016-11-30
DLP-847	Meet camera DAQ/OCS milestones for FY16 *This is one of the delayed milestones reported at the July 2017 DM Review. It is expected to be completed during the the F17 cycle. There is no impact on the critical path.*	2016-11-30
	International Communications & Base Site	
DMTC-7400-2420	REUNA La Serena - Santiago DWDM equipment *Completion was delayed by 1 month by late delivery of equipment from vendor (Coriant). No impact to critical path*	2017-07-31
DMTC-7400-2090	Report on functional fiber connections, including AURA equipment *Completion was delayed by 1 month due to delayed access to the Cerro Pachon caseta. A new backup power generator is required and AURA Operations is in the process of procuring and installing the generator. No impact to critical path.*	2017-08-18

Science Quality & Reliability Engineering		
DLP-579	Usability and deverloper efficiency *Will be completed later in F17.*	2016-08-31

Detailed Project Progress

1.02C.01: System Management

Current accomplishments

The DM Project Manager:

- Presented DM and especially the Science Platform to the NOAO LSST Community steering committee in Tucson Aug 7th.
- Attended the LSST Project and Community workshop. Specifically, he helped organise the Commissioning Rehearsals which were quite informative - we stepped through 4 points in 2 days of commissioning to see where people thought they would be physically and which tools they were using and to see the timezones and what they mean.
- Prepared for the Joint Status Review. The Deputy DMPM (Swinbank) and Jeff Kantor successfully took DM through the Director's Review Aug 23 and 24. The fact that such a review can be delegated shows a good level of stability in DM.
- Attended “The Science of Gaia and Future Challenges” - this was very informative and many useful slides ets were shared with the DM Science Team and beyond. Spoke to several contacts about drumming up candidates for DM Subsystem Scientist role.

Planned activities

The DM Project Manager will:

- Get DM through the Joint Status Review.
- Present DM at AMCL.

1.02C.02.01: Data Management Science

Current accomplishments

The DM Subsystem science team focused on completing materials for and participating in the LSST2017 meeting (Juric: LSST Science Platform, Graham: LSST Special Programs), including a number of presentations on DM deliverables to the advisory bodies and the

broader community.

Melissa Graham prepared for and attended the LSST Community Science Center working group's face to face meeting at NOAO, where she gave a presentation on LSST data products, and contributed to writing the white paper report. Completed a revised draft of a DM technical note reviewing and proposing changes to the Data Management plans for processing of data from Special Programs.

Slater presented a draft of the database performance requirements document to the database and architecture teams at LSST2017. Developed some potential methods for interfacing the DRP database with map-reduce type software, which he began to prototype during the rest of the month. Continues to supervise Suberlak on crowded fields.

Suberlak deployed the LSST Stack photometry tools on DECam data for fields representative of selected source density levels, and is working on comparing the results to TRILEGAL and *galfast* simulations.

Juric strategically positioned the DM-community communication plan by socializing it within the Project. With Colin Slater, he held a tiger-team lock-up session to understand the science requirements for data access. He prepared the briefing on the LSST Science Platform for the Science Advisory Committee, which has been postponed for September. Majority of the remaining effort was spent understanding the best big-data management software options to use, in conjunction with Level 3 processing schemes. The planned activities were within 10% of the goal.

Planned activities

- Shepherd to adoption outstanding LCRs and RFCs whose resolution was agreed to in principle at LSST2017.
- Continue the work on photo-z, both addressing the referee report from the submitted paper, and working on the next one which incorporates IR.
- Advance the science requirements for the database.
- Advance the crowded field study.
- Juric will be working on winding down DM Subsystem Science activities and transitioning into the DM SST Coordinator role.

1.02C.02.02: DM System Architecture

Current accomplishments

The DM Software Architect:

- Discussed Alert Distribution design; wrote up early draft document and diagram; iterated once.
- Attended LSST2017 PCW; clarified uses and discussed options for non-database

catalog processing, workflow, Data Butler, Camera Diagnostic Cluster, Auxiliary Telescope integration.
- Discussed Calibration Database, provenance database and capture, and L1 database prototype implementation with DAX team.
- Reviewed "recipe" configuration for serializers.
- Prepared for study of cloud deployment options.
- Discussed Integrated System Parameter Database with Systems Engineering personnel; started preparing a list of relevant parameters.
- Participated in Unit/DataId/Registry design discussion.
- Reviewed LSE-309.
- Responded to question from Directors' Review.
- Prepared for Annual Review.

The DM Systems Engineer:

- Completed the LSE-61 verification matrix.
- Worked on the LSE-131 verification matrix.
- Attended LSST2017 meeting and aided with Python 3 porting hack session.
- Modified sconsUtils to run pytest for test execution and worked with SQuaRE to integrate into Jenkins CI. We have now switched over completely to pytest and use multi-process testing.
- Convened the Butler Working Group and had multiple meetings where we have discussed common butler nomenclature and scope and written many use cases.
- LCR-921 ("DM to Provide Service for Publishing Predicted Visit Sequence") was approved and a new version of LSE-61 released.
- LCR-1024 (“Tie LSE-163 into OSS”) was voted upon and approved.
- LCR-1064 (“Add priorities to LSE-61”) was submitted.

The DM System Interfaces Scientist:

- Prepared an updated and more detailed proposal for cross-subsystem early integration activities.
- Attended the 2017 LSST PCW. Presented the early integration schedule at several sessions. Discussed the early-integration schedule and the auxiliary-telescope spectrometer deployment schedule. Prepared a revision to the early integration schedule to reflect updates from the aux-tel session.
- Presented material on the Portal Aspect of the Science Platform at a PCW session on “what to expect at a DAC”.
- Attended the inaugural “JupyterCon” conference on the “Jupyter ecosystem” - the open-source software that LSST will use for the Notebook Aspect of the Science Platform. Prepared detailed notes on sessions attended. Started discussions with members of the Jupyter and Cyverse teams on future collaboration.
- Provided the updated early-integration schedule to the SE team for incorporation in the Status Review materials and PMCS.

Planned activities

The DM Software Architect will:

- Update the Data Backbone descriptions in LDM-148 and LDM-152.
- Plan a design review for the Header Service.
- Work on hiring a new Architecture Team member.
- Present as needed at the Joint Status Review.
- Continue to prepare for cloud deployment study; help prepare for SQL Server study.
- Take two weeks' vacation.

The DM Systems Engineer will:

- Continue with the Butler Working Group to write a requirements document.
- Investigate adding code coverage to the test suite.
- Continue to work on requirements in LSE-61 and DM subsystems.

The DM System Interfaces Scientist will:

- Work on clearing out old Systems Engineering action items ("LIT"s) relevant to DM and on clarifying which ones need substantive work from others in DM.
- Resume work on LSE-75 revision.

1.02C.03: Alert Production

Current accomplishments

02C.03.00 – Alert Production Management Engineering and Integration

- Travel & meetings:
    - The team spent the week of 14 August at the LSST2017 Project & Community Workshop in Tucson, AZ.
    - Swinbank spent the week of 28 August visiting the team at UW and discussing ongoing work in preparation for his assuming the role of T/CAM.
- Construction of an "end-to-end" Alert Production test system [DM-9676, DM-10770 DM-10771, DM-10773, DM-10775]:
    - Development of the end-to-end AP test system has been the major focus of the group through this month. The aim of this system is to demonstrate alert processing from science data through to detection of transient and variable objects in a minimally functional way; it will then be extended and refined to meet all of LSST's science requirements.
    - The group organized a session discussing this system at the LSST2017 Project & Community Workshop. [DM-11377]
    - A collection of HiTS data, recorded using DECam, has been selected as the

primary data source for the initial end-to-end system. This data has now been collected, packaged appropriately, and made available on the LSST development cluster at NCSA. [DM-10976, DM-11371, DM-11575]
    - Work has continued in developing the pipeline script which is used to execute the AP pipeline. This included making it more generic, so that it relies on fewer arbitrary hard-coded variables, and updating its configuration and layout to be compatible with the standard conventions used within the DM Stack. [DM-10975, DM-11021, DM-11324, DM-11325, DM-11461]
    - A key part of this work is not simply to demonstate the successful execution of the pipeline, but to record "metrics" describing its behaviour and which can be used to characterize and diagnose performance issues. The first — and simplest — set of metrics, describing the run-time performance of the system, have now been developed. [DM-10779]

02C.03.01 – Single Frame Processing Pipelines

- No work was undertaken in this WBS during this month.

02C.03.02 – Association Pipelines

- Protoype source association system [DM-10768]:
    - Source association — gathering together source detections in individual epochs to create distinct astronomical objects — is an essential component of the end-to-end system (described above). A basic association system is therefore under construction to support that activity. This system will be extended to support more advanced association algorithms later in construction.
    - The association system depends on building collections of objects detected on difference images. Work in this month focused on developing the API necessary to work with those collections. [DM-11283]

02C.03.03 – Alert Generation Pipelines

- Testing prototype alert distribution system [DM-7355]:
    - Tools for monitoring the status and throughput of the prototype alert distribution system have been developed. [DM-8820]

02C.03.04 – Image Differencing Pipeline

- Include decorrelated Zackay et al. (2016) and Alard & Lupton (1998) image differencing algorithms in the LSST stack [DM-9989]:
    - The DM Stack's image differencing task has been updated to include a spatially-variable ZOGY image differencing algorithm. [DM-10805]
- Incorporate Differential Chromatic Refraction (DCR) in the DM Stack [DM-9613]:
    - Simulated images with appropriate metadata for testing the DCR correction

algorithm have been generated. [DM-9900]

02C.03.05 – Application Framework for Exposures

- Emergent work [DM-9680, DM-10068, DM-10774]:
  - These activities includes essential maintenance and general infrastructure upgrades.
  - Work focused on upgrading specific packages to use the Pytest unit testing infrastrastructure, an activity being coordinated by the System Architecture Team. [DM-11594, DM-11585, DM-11429, DM-8688].
  - LSST contributed code to the external AST package to correctly round-trip zero-length vectors through coordinate transformations. [DM-11454]
  - The test suite for camera distortion was upgraded. [DM-11574]
  - LDM-294, the DM Organization and Management description document, was updated to include the new WBS which will be used for the 02C.03 group when the DM replan has been baselined.
  - This completed activity DM-9680.

02C.03.06 – Moving Objects Pipeline

- No work was undertaken in this WBS during this month.

02C.03.07 – Photometric Calibration Pipeline

- Work continues on developing the Jointcal system in preparation for comparative testing against the older meas_mosaic system. This work was not concluded during this month. We expect testing to start taking place during September.

02C.03.08 – Astrometric Calibration Pipeline

- No work was undertaken in this WBS during this month.

Planned activities

02C.03.00 – Alert Production Management Engineering and Integration

- Work will continue on the development of the end-to-end testing system.

02C.03.01 – Single Frame Processing Pipelines

- No work is planned in this WBS during this month.

02C.03.02 – Association Pipelines

- Work will continue on the association system in support of end-to-end testing.

02C.03.03 – Alert Generation Pipelines

- Benchmarking and monitoring of the alert generation system will continue.

02C.03.04 – Image Differencing Pipeline

- Integration of the decorrelated Zackay et al. and Alard & Lupton differencing algorithms into the DM Stack will be concluded, and the process will be documented, ahead of the the departure of David Reiss who has spearheaded this activity but is now leaving the project.
- Work will continue on developing and documenting the DCR correction capabilities of the stack.

02C.03.05 – Application Framework for Exposures

- We will continue to service technical dept and emergent requirements.

02C.03.06 – Moving Objects Pipeline

- No work is planned in this WBS during this month.

02C.03.07 – Photometric Calibration Pipeline

- Comparison of Jointcal with meas_mosaic will begin. Ultimately, this should lead to the latter being retired.

02C.03.08 – Astrometric Calibration Pipeline

- No work is planned in this WBS during this month.

Recruitment update

- David Reiss will leave the AP team at the end of September. A search is underway for his successor.
- Simon Krughoff has stepped down as T/CAM for the AP team; he will be replaced by John Swinbank.

1.02C.04: Data Release Production

Current accomplishments

02C.04.00 – Data Release Production Management Engineering and Integration

- Travel & meetings:
    - The DRP team attended the LSST2017 Project & Community Workshop in Tucson, AZ for the week 2017-08-14.

    - Swinbank visited the University of Washington, Seattle, WA for the week of 28 August, to work with the team there in preparation for taking up his new role as T/CAM.
    - Swinbank participated in the LSST Directors Review during the week of 21 August.
- Management & planning [DM-10816]:
    - LDM-564, the DM release description document, was updated to include clearer descriptions of the Science Pipelines functionality which will be available in each major release of the DM Stack. [DM-11517]
    - In response to feedback from the July 2017 review of DM, an audit was carried out of all milestones from the previous year of development which had not been accomplished. Some of these are late; some are obsolete because of the DM replan; some have been accomplished, but not properly recorded as such. Status was recorded for each and this was presented to reviewers during the Directors Review. [DM-11533, DM-11534]
- Forward Global Calibration Method (FGCM; Burke et al., 2018) [DM-10584]:
    - Work has continued on incorporating the FGCM into the DM Stack. This has not yet been completed. [DM-10440]
- Emergent work and pipeline support [DM-10382]:
    - As part of the System Architecture coordinated effort to update our unit test system to use the Pytest framework a number of numerical issues caused by (mis-)use of global state in our existing tests were discovered. These have been resolved. This has both enabled the upgrade to Pytest and increased the numerical rigour of our test suite. [DM-11620]
    - More flexibility has been added to the selection of the statistical operation which is used to combine data to made coadds. In particular, it is now possible to generate coadds using the median, but use of other operations is also possible. [DM-4158]
- Pipeline QA [DM-10571]:
    - A number of upgrades have been made to the QA “dashboard” system which is under development. In particular, these have included an investigation into possibilities for integrating the system with JupyterLab-based notebooks and upgrading the system to deal with much larger data volumes. [DM-11254, DM-11255, DM-11272, DM-11405]

02C.04.01 – Application Framework for Catalogs

- Middleware and framework development [DM-10586]:
    - The DRP team has contributed to the final stages of the SuperTask Working Group, writing documentation and checking requirements. [DM-10845]
    - The DRP team is also actively participating in the Butler Working Group, contributing use cases and developing design sketches of the data repository. [DM-11389, DM-11573, DM-11721]

- Emergent work and reduction of technical debt [DM-10383]:
    - A number of minor build and test problems were resolved across the DM Stack. [DM-7457, DM-10182, DM-11158, DM-11507, DM-11539]
    - A numerical issue in the core measurement algorithm package, meas_base, has been resolved. This clears the way for DM to update to a more recent release of the GCC compiler system in future. [DM-10902]

02C.04.02 – Calibration Products Pipeline

- Auxiliary telescope development [DM-10581]:
    - No significant progress has been made during this period due to unavailability of key personnel.
- Processing camera test stand data [DM-10897]:
    - obs_comcam, which provides integration between test stand (and ComCam) data and the LSST stack, has been validated and made available within the DM Stack. [DM-9872]
    - Work is underway to build the first version of the calibration products pipeline, which will calculate per-amplifier gains based on 55Fe data. [DM-11479].

02C.04.03 – PSF Estimation

- Wavefront measurement and PSF reconstruction [DM-10355]:
    - The work required to develop a reference wavefront in the unrotated frame, described in last month's report, has now been completed successfully and merged to the DM Stack. [DM-10850]
    - Work continues on a technical note describing the worried carried out to date in this area. [DM-10355]

02C.04.04 – Image Coaddition Pipeline

- Warped image comparison [DM-8290]:
    - Work has continued on refining and implementing the "Robust Coaddition" described last month into the DM Stack, where it is now referred to as CompareWarpAssembleCoaddTask. This work had been broadly completed at the end of August, but had not yet been reviewed and merged to the Stack. At the same time, scientific tests of the new algorithm are ongoing. [DM-11432, DM-11445]
    - Work has started on refactoring the coaddition tasks to a use new, more flexible, way of referring to different types of coadd. [DM-11446]

02C.04.05 – Object Detection and Deblending

- Deblender development [DM-10353]:
    - An investigation into using source colors to estimate the bounding box of

each source has been completed successfully. The feasibility of such an approach has been demonstrated. However, developing this into an effective algorithm which can be merged to the DM Stack will require substantially more work. DM-11720
    - The prototype deblender code now allows sources to be “translated”, or shifted, by a small distance relative to the initial estimate of their position in order to improve deblending results. [DM-10189]

02C.04.06 – Object Characterization Pipeline

- Experiments in shear measurement on coadds [DM-10579]:
    - The Bayesian Fourier Domain (BFD) codebase has been updated to correctly account for correlated noise. [DM-10824]
    - Initial experiments in shear measurement have been completed. The scope of this work was included to include covariance since DM-10824 was successfully completed first (see above). The BFD and Metacalibration methods were in agreement and did not show a bias on the shear. [DM-10825]
- Improved galaxy flux measurement algorithms [DM-10580]:
    - Due to staff parental leave, no work was undertaken on this activity during August.

Planned activities

02C.04.00 – Data Release Production Management Engineering and Integration

- Swinbank will Attend the LSST Joint Status review in Tucson during the week of 2017-09-04.

02C.04.01 – Application Framework for Catalogs

- Continue to participate in the Butler Working Group and ensure all SuperTask WG design documents are completed.

02C.04.02 – Calibration Products Pipeline

- Analysis of test stand data will continue.

02C.04.03 – PSF Estimation

- The DM Technical Note describing work carried out in this WBS will be published.

02C.04.04 – Image Coaddition Pipeline

- Complete the merger of CompareWarpAssembleCoaddTask to master.

02C.04.05 – Object Detection and Deblending

- Work towards at-scale testing of the deblender prototype.

02C.04.06 – Object Characterization Pipeline

- Continue to compare techniques for shear measurement on coadds.

Recruitment update

- Sophie Reed will join the group as a postdoctoral scientist in early September.
- John Swinbank will move to the University of Washington at the end of September, but will continue to serve as T/CAM for DRP.
- Advertisements have been issued for a new project manager, a software developer, and multiple software scientists. Applications will be reviewed over the next several months.

1.02C.05: Science User Interface & Tools

Current accomplishments

02C.05.00

- Continued to work with IPAC IRSA group on collaboration in Firefly development, plan and schedule coordination.
- All attended LSST2017 community workshop and all hands meeting.
- Coordinated the schedule for kick-off meeting of "LSST authorization authentication and identity management."
- Summer vacation: 4 members, total 23 days.
- Gregory attended JupyterCon at NY.

02C.05.01 Basic Archive Access Tools

- Many bug fixes in Firefly reported by testers in IPAC IRSA group.
- FIrefly was deployed as a Docker container at NCSA. It is able to access GPFS file system , making it possible to browse through re-processed HSC images at NCSA.
- Outline for Firefly testing procedure.
- Manual UI testing.

02C.05.02 Data Analysis and Visualization Tools

- Firefly_client
    - Improvement to Firefly Python API (DM-11606).
    - Demo to use the new grid layout (DM-11476).

02C.05.03 Alert/Notification Toolkit

- No new work done.

02C.05.05 User Workspace

- No new work done.

Planned activities

02C.05.00

- Continue to work with IPAC IRSA group on collaboration in Firefly development, plan and schedule coordination.
- Support for joint status review.

02C.05.01 Basic Archive Access Tools

- Start tackling some UX issues.
- Work on security and authorization related issues.
- PDAC development:
    - Work with metaserv v1 API (DM-8215).
    - Prep-work for NEOWISE data.

02C.05.02 Data Analysis and Visualization Tools

- New functions
    - Support image drawing layer sub-grouping
    - Make plot.ly plotting api to handle unrecognized chart types cleanly (DM-11477).
- Improvement for plotting API.
- Finish Firefly_client Python documents and examples.

02C.05.03 Alert/Notification Toolkit

- No work planned.

02C.05.05 User Workspace

- No work planned.

Recruitment update

- No activities.

1.02C.06: Science Data Archive & Application Services

Current accomplishments

02C.06.00 Science Data Archive and Application Services Management Engineering and Integration

- Team continued improving robustness of DAX services and otherwise supported SUIT integration efforts in the PDAC as necessary.
- Entire team attended 2017 LSST Project and Community Workshop in Tucson, AZ.
- Gates, Gaponenko, Hanushevsky, Pease, Jammes, and Mueller had vacation during this month.

02C.06.01.01 Catalogs, Alerts and Metadata

- No work was carried out in this WBS element this month.

02C.06.01.02 Image and File Archive

- Lo implemented a command-line interface for the ImgServ web service, and an associated set of unit tests for the service. [DM-10807, DM-10909]
- Lo implemented various low level support features within ImgServ in support of upcoming SIAv2/SODA cutout interfaces. [DM-10427, DM-10450]

02C.06.02.01 Data Access Client Framework

- Pease addressed obs_test teardown race conditions that were causing intermittent CI failures. [DM-11735]
- Pease addressed a CI test failure in obs_sdss. [DM-11691]
- Pease cleaned up an unnecessary dependency of the templates package on daf.fmt.mysql. [DM-11494]
- Pease added error detection/reporting code for empty mapper location templates in dataset definitions. [DM-8432]
- Pease added error detection/reporting for non-dict RepositoryArgs.mapperArgs. [DM-11289]

02C.06.02.02 Web Services

- Van Klaveren implemented Pandas-style to_sql and from_sql for afw.table via SQLAlchemy. [DM-10376]

02C.06.02.03 Query Services

- Salnikov implemented result retrieval for Qserv asynchronous/disconnected queries. [DM-11191]
- Salnikov completed port of Qserv administration scripts and ancillary tooling to

Python 3. [DM-11550]
- Salnikov ported C++ Python wrappers in Qserv from SWIG to pybind11. [DM-11593]
- Salnikov ported C++ Lua wrappers in Qserv from SWIG to native C++, removing the final dependency on SWIG within the stack. [DM-11599]
- Mueller addressed a faulty unit test for the Qserv qhttp modules that was causing intermittent CI failures. [DM-11487, DM-11615]
- Mueller updated scisql package to patched upstream, and addressed a few miscellaneous compiler warns. [DM-11608, DM-11538]
- Gates addressed memory over-allocation issue that was impacting Qserv czar performance while aggregating large query results. [DM-10566]
- Gates performed network and node failure tests on query cancellation cleanup feature, and implemented some fixes and performance improvements based on the test results. [DM-11172, DM-11524]

02C.06.02.04 Image Services

- No work was carried out in this WBS element this month.

02C.06.02.05 Catalog Services

- No work was carried out in this WBS element this month.

Planned activities

02C.06.00 Science Data Archive and Application Services Management Engineering and Integration

- Team to continue improving robustness of DAX services and otherwise support SUIT integration efforts in the PDAC as necessary.
- Mueller, Pease, and Hanushevsky have vacations scheduled this month.

02C.06.01.01 Catalogs, Alerts and Metadata

- No work is planned for this WBS element this month.

02C.06.01.02 Image and File Archive

- Lo to complete initial implementation of image service supporting image cutouts via IVOA SODA protocol.

02C.06.02.01 Data Access Client Framework

- Pease to implement afw BaseCatalog reader and writer for Butler.
- Salnikov to investigate Bosch's latest SuperTask designs.

02C.06.02.02 Web Services

- Van Klaveren to begin working on support for asynchronous queries at the web

service layer.

02C.06.02.03 Query Services

- Hanushevsky to complete converting czar code to use the new XRootD SSI v2 API for content-addressed messaging between czar and workers.
- Jammes to complete refactoring MariaDB into an independent container in Kubernetes-based Qserv deployments.
- Gaponenko to continue work on Qserv data distribution/replication framework.
- Gates to tie up implementation loose ends with Qserv query cancellation and large-result handling.
- Salnikov to implement status monitoring and cancellation for Qserv asynchronous/disconnected queries.
- Jammes to continue adapting Qserv for deployment within Kubernetes.

02C.06.02.04 Image Services

- No work is planned for this WBS element this month.

02C.06.02.05 Catalog Services

- No work is planned for this WBS element this month.

Recruitment update

- No recruitment activity this month.

1.02C.07: LSST Data Facility

Current accomplishments

02C.07.00 – Processing Control and Site Infrastructure Management, Engineering and Integration

Planning and Management

These activities involve the creation of project plans for the construction phase for the DM-wide replan, high-level engineering and design specification, and planning for Data Facility operations during the commissioning and operations phases. NCSA leads planning for operations of the Data Products Production (DPP) department.

Following the completion of the DM replan, the goal of the LSST Data Facility management is to render the architecture and plan in a sustainable way both for internal management and interfacing with the overall project. We began an investigation of the TOGAF standard enterprise architecture language ArchiMate, and associated open source tool Archi, for documenting and managing the logical architecture of the LSST Data Facility services and dependencies. This provides a layered, coherent model of the facility, with viewpoints

supporting management, engineering, and planning, and a configuration management database at the service level. We began prototyping with diagrams of Level 1 archiving services, consolidated database services, and site and Data Backbone file systems. We discussed our modeling framework with the LSST Systems Engineering and DM Architecture teams at the LSST all-hands meeting about interfacing with the SysML-based Model Based System Engineering (MBSE) models.

Several NCSA staff prepared for and participated in the LSST2017 Project and Community Workshop, including presenting current services the Data Facility provides to the project, presenting the Data Facility roadmap, presenting on recent work in setting up batch production services, planning for commissioning, and reviewing data retention and access policies.

We made a preliminary definition and costing of descope options for the Data Management replan; see DM-11466. We reviewed FY18 priorities and milestones with technical area leads and specified use cases for the FY18 acquisition strategy plan.

- DM-10692 Overall Facility Planning and High-Level Engineering

Service Management and Monitoring

The Service Manager develops and maintains the service catalog, including service-level targets, service-level agreements, and principal technical and managerial liaisons. This includes monitoring and managing availability, capacity, and IT continuity of services, and other operational matters, as well as forming a definitive opinion about the satisfaction of each customer for the services provided. The service manager supports deployment of services to operation and service transitions.

We finished development of an initial request and incident receipt process, along with JIRA submission forms providing a single point of contact for users and developers to report issues and make requests. This first version is facilitated by an internal process classifying tickets, assigning them to staff, and tracking the resolution time. Improvements to both the processes and the systems to support service management functions are needed and will be a focus of the next cycle. For service-level monitoring, we developed a prototype quad chart display summarizing system-level Nagios data and other sources. In emergent work, we worked with Qserv developers to complete specification of an expansion of development hardware.

- DM-9658 Enterprise Monitoring Design - Phase 2
- DM-10693 Service Management Service Operation Workflow Development
- DM-10695 Service Management - Emergent (F17a)

Hardware acquisition and provisioning planning

These activities involve high-level planning and design of physical systems related to computation, storage, networking, administration, and IT security on which DPP services

are running or will be deployed. This includes consideration of near-term construction phase needs, as well as preparation for commissioning and operational needs.

We finalized planning for provisioning the full-scale L1 Complete Test Stand, a refresh of the lsst-db system, the raft-scale Oracle database system, and a storage capacity increase. Procurement activities for these capabilities are being drafted. We began requirements collection and sizing determination for FY18 acquisition planning.

- DM-10696 FY17 Procurement Activity #2 Planning
- DM-10697 FY17 Procurement Activity #3 Planning
- DM-10698 FY18 Acquisition Planning

02C.07.01 – Processing Control

Batch Production Services

These activities support the construction of services that enable the production of data products in a batch environment (e.g., Level 2 data products). The batch production services will execute processing campaigns on computing resources to produce the desired LSST data products, where campaigns are defined as sets of pipelines (ordered ensembles of computational steps), inputs they are being run against, and methods handling their outputs.

We continued development of the offline Batch Production System, demonstrating an execution of production workflows, focusing on handling large workflows, dealing with various failure scenarios, and assessing Pegasus functionality and performance. We continued biweekly reprocessing of the HSC RC dataset based on DM needs, to both verify the software stack and to test the development of science algorithms. We also finalized operational analysis of the initial HSC full reprocessing campaign that took place in May, including summarizing operational metrics such as capacity usage, processing times, and problem management (root-cause analysis).

- DM-10700 Batch Processing System Phase 1b (Pegasus)
- DM-10702 HSC reprocessing campaigns

Level 1 Services

These activities support the construction of services that will support Observatory Operations and enable the production, transport, and archiving of Level 1 data, including the nightly stream of images and events collected and processed in near real time.

We continued refining the front-end interfacing elements of the archiving and prompt processing Level 1 services, designing and implementing a prototype for science cluster needs, writing specifications for an auditing system, and continuing coding on the interface to the DAQ using the system installed in the Level 1 Complete Test Stand in July. We continued work on the header service implementation to integrate EFD and OCS metadata,

completing the implementation of python multiprocessing in the header writer. The header service design was presented to stakeholders at the LSST2017 workshop.

- DM-10703 L1 Services & Integration Activity Support
- DM-10704 Header Service Design & Implementation

Common Middleware and Other Tasks

These activities consist of developing and maintaining general processing control software components, including common software used in production processing and by DM developers, as well as emergent unplanned support tasks.

Emergent work for August included involvement with the Butler Working Group—attending meetings and providing input into use cases from a production operations viewpoint—as well as working through issues arising during HSC reprocessing.

- DM-10705 Emergent Middleware Work (F17a)

02C.07.02 – Infrastructure Services

Data Backbone

These activities support the construction of services which archive, catalog, and distribute data to compute resources and data access endpoints across all sites—ingestion, file management, database hosting, etc.

We continued investigating methods for data movement over a WAN to support the spectrograph test stand in FY18, as well as long-term solutions for data distribution in the Data Backbone, completing a first phase investigation of the Rucio data distribution package. We set up a storage element with a frontend gridFTP server and used Rucio to upload data into the storage element, recording files into the Rucio file catalog. This demonstrates the ability to move data from a remote site to NCSA.

Work on the consolidated database services included setting up the prototype Oracle system and supporting its use in testing the Dark Energy Survey Data Management) DESDM framework for LSST use.

- DM-10707 File distribution and management framework
- DM-10709 Operational concepts of databases supporting production data flows

Other Infrastructure

These activities support the construction of general reusable services and infrastructure components, as well as emergent infrastructure work not included above.

We continued working with Kubernetes, HTCondor, and Docker on a small dedicated computing cluster to prototype on-demand, elastic provisioning and understand resource management capabilities. Investigations of configuring and managing systems for

containerized applications continued, including testing authorization and authentication for users and beginning to evaluate the use of OpenShift for managing Kubernetes containers. Requirements discussions with DM developers constructing containerized applications continued.

- DM-10711 Containerized Application Management Services

02C.07.03 – Environment and Tools

No work is planned in this WBS for the F17 cycle.

02C.07.04 – Site Infrastructure

Hardware Acquisition, Deployment and Provisioning

These activities consist of the acquisition, provisioning, configuration, maintenance, and decommissioning of physical capabilities and associated systems which support LSST services running at NCSA.

DAQ test stand deployment activity was completed with security vetting and follow-up. We continued receiving components for the Chilean Base instance of network-based security (intrusion detection, log collection, scanning, etc.) authentication & authorization systems, setting up a test rack to accommodate verification and burn-in. NCSA 3003 lab refresh activities continued, with full deployment of the reconfigured racks and power infrastructure; minimal cleanup and decommissioning remains.

- DM-8510 DAQ Teststand Single Raft Deployment
- DM-10712 Chile Base AA Acquisition & Installation
- DM-10715 NCSA 3003 Lab Refresh

Service Management for LSST developer resources, Emergent and Miscellaneous Work

These activities consist of supporting and communicating with users of the current services provided by NCSA to LSST DM developers, as well as general improvements and additions to hardware and system infrastructure.

Work on Disaster Recovery (DR) design and implementation completed, with verification of successful backup and restore of the /datasets, /home, and /software partitions, as well as security vetting and improvements to process automation scripts. The current phase of system monitoring improvement was completed, with initial monitoring of the PDAC and Nebula systems based on DM input, as well as improvements and extensions to configuration (puppet), low-level system parameter data collection, and monitoring of data collection. Work continued responding to the needs of developers deploying a prototype Elastic Services cluster.

- DM-8502 Disaster Recovery for Science Datasets
- DM-10716 System Monitoring Refinement (F17a)

- DM-10719 Elastic Services Testbed Provisioning

Planned activities

02C.07.00 – Processing Control and Site Infrastructure Management, Engineering and Integration

Planning and Management

In September, we will begin working on articulating requirements documents and defining verification tests consistent with the Data Facility and overall DM roadmap, and begin breaking down replan planning packages into work packages for the next planning cycle.

We will continue work articulating the Data Facility roadmap, milestone matrix, and service dependencies in a configuration management database (CMDB). The goal is to record the service dependencies, components, and processes, which supports overall management of the facility, provides a basis for detailed cycle planning, and supports change control of Data Facility services. We continue use of the TOGAF-standard tool, Archimate, to model the CMDB. Priorities include documenting the Data Backbone set of services and the DACs.

We will continue to participate in the DM subsystem science team and DM subsystem engineering team and associated activities, including leading the compression working group to confirm the factor of FITS compression that the Dark Energy Survey has demonstrated is scientifically viable for LSST, and costing the proposed data retention plan of storing processed visit images (PVIs) on disk. We will also participate in the NSF/DOE Joint Status Review as needed.

Finally, in September we will prepare the FY18 contracts and labor and non-labor budgets.

- DM-7632 Develop planning packages from DM replanning deliverables
- DM-11836 Facility Planning & Management (F17b)

Service Management and Monitoring

We will continue maintaining and evolving the Service Management framework, and refining the Change Control and Incident Response frameworks. Differentiation of tickets according to incidents, requests, and changes, as well as the processes for handling each of these functions, is the next objective of development. We will continue to work toward integrating services into the service-level monitoring framework to support FY18 project milestones.

- DM-10693 Service Management Service Operation Workflow Development
- DM-10694 Service Management Framework Implementation
- DM-10695 Service Management - Emergent (F17b)

Hardware acquisition and provisioning planning

We will continue planning for procurement of the remaining FY17 hardware expansion as described in the FY17 Annual Acquisition Strategy Document. Procurements include additional infrastructure for the network security deployments in Chile, the raft-scale Oracle database, and the GPFS storage expansion. We will continue planning FY18 capabilities and prepare the FY18 Acquisition Strategy document.

- DM-10697 FY17 Procurement Activity #3 Planning
- DM-10698 FY18 Acquisition Planning

02C.07.01 – Processing Control

Batch Production Services

We will continue to enhance the DESDM-based processing service based on the need for a stable production-like system in the near term before commissioning ramps up. We will continue the planned reprocessing of HSC data based on DM needs, to both verify the software stack and to test the development of science algorithms. The next HSC reprocessing campaign is targeted to begin in October to meet the DM milestone in November, and preparation includes scaled testing using the DESDM framework and the Oracle test system. We will also document our investigation of Pegasus failure cases based on operational use cases.

- DM-10701 Batch Processing System Phase 1a (DES)
- DM-11839 Dataset Reprocessing
- DM-11840 Batch Production Development (F17b)

Level 1 Services

We will continue refining the archiving and prompt processing Level 1 services, using the established Early Pathfinder Software Integration Activity schedule to drive development priorities. Specific activities in September include demonstrating readout of the Camera DAQ system and preparing for the October early pathfinder activity which demonstrates response of the OCS “next visit” event. We will also continue working with Camera and Telescope & Site developers to identify sources and contents of the EFD, OCS metadata available to build headers, and contents of headers required by CCS developers and downstream DM processing. Documentation of the header service will be drafted. We will continue working with the installed DAQ simulator, integrating and testing our framework with the provided API. Upgrades to the archiver package to support its use on the test stand in Tucson will be built and tested in this environment.

- DM-11842 L1 Software Improvements phase 1
- DM-11841 DAQ Software Support phase 1
- DM-11838 Header Service

02C.07.02 – Infrastructure Services

Data Backbone

We will continue investigating methods for data movement over a WAN supporting the spectrograph test stand in January 2018, as well as long-term solutions for data distribution in the Data Backbone. In September will investigate how to integrate the Rucio framework and catalog with production file management of the central archive.

We will continue participation in the Butler working group, representing Data Facility operations, central archive management and data distribution, and production requirements.

- DM-11844 File distribution & Management (F17b)

Other Infrastructure

Work continues on investigating and prototyping the containerized application management framework, including the interaction between Kubernetes, Docker, and HTCondor, as well as the resource management capabilities of elastic provisioning. We will continue discussions with relevant parties, e.g., the HTCondor team, Pegasus team, and SQuaRE (integrating Jupyter notebook containers). We will also review security concerns related to internal use of elastic services by LSST staff and use by external users.

- DM-11843 Containerization phase 2 (F17b)

02C.07.03 – Environment and Tools

No work is planned in this WBS for the F17 cycle.

02C.07.04 – Site Infrastructure

Hardware Acquisition, Deployment and Provisioning

The Chile Base instance of network-based security and authentication & authorization systems will be provisioned locally, configured, and verified, in preparation for shipment and installation in Chile. Decommissioning and cleanup for NCSA 3003 lab refresh activities will be completed. Acquisition and provisioning of the remaining FY17 hardware procurements will continue as these final FY17 procurements are completed.

- DM-10712 Chile Base AA Acquisition & Installation
- DM-10715 NCSA 3003 Lab Refresh
- DM-10719 Elastic Services Testbed Provisioning
- DM-10714 Level 1 Test System Full Provisioning
- DM-10717 Prototype Consolidated Database Provisioning
- DM-10718 Database (lsst-db) Refresh and Capacity Increase Provisioning
- DM-10722 GPFS Storage Expansion Provisioning

Service Management for LSST development resources, Emergent and Miscellaneous Work

Work will continue on system monitoring improvements, with the goal of providing high-level service and system monitoring displays to support FY18 project integration activities, and supporting development of service-level monitoring and management framework and process development. We will also continue normalizing administration and configuration of hardware systems across enclaves and sites.

- DM-11845 System Monitoring Refinement (F17b)
- DM-10725 Consolidation of ITC Management Systems (F17a)

Recruitment update

A requisition for a full-time database administrator is in progress.

1.02C.08: International Communications & Base Site

Current accomplishments

02C.08.01 – Base Center

- Summit Base ITC Tiger Team: LCR-964 was approved to baseline LSE-299 Summit Computer Room Requirements Document.
- Summit and Base Networking and Computing: Ron Lambert and Luis Corral spent at the LSST2017 meeting, followed by a week in Tucson. In this week, they were at Cisco's R&D site conducting a Cisco Proof of Concept and training on the ACI network. Also while in Tucson, Ron Lambert and Luis Corral talked to Bill Schoening with regard to cable ducts in the Telescope. Bill will work with them in preparing diagrams to deliver to a contractor. NCSA was informed that the La Serena - Santiago network link would not be ready for their Identity Management system/network installation due to delay in receiving REUNA DWDM equipment from Coriant. There is no point in having NCSA come down to install the equipment before the network link is operating. We expect this will be in November.

02C.08.02 – Chilean Data Access Center

- No activity this month.

02C.08.03 – Long-Haul Networks

02C.08.03.01 – Chile National WAN

- Summits–AURA Gatehouse Network: No activity. This segment is accepted.
- DWDM Equipment: There were delays in the receipt and installation of the equipment. While this will not impact LSST commissioning, it will delay the First Optic Light Demonstration into November. The AURA DWDM installation has been

delayed due the Caseta on Cerro Pachon not having enough backup generator capacity per AURA Operations. AURA Operations has started the process to procure and install a new generator, but we expect at least a one month slip in the AURA DWDM installation. The REUNA DWDM equipment was not received from Coriant on the original schedule. Current expected dates for the receipt are:
    - Intermediate nodes (amplifier nodes): 30/Aug arrived om Santiago.
    - Terminal nodes (chassis of La Serena – Santiago backbone and Santiago metropolitan ring): 8/Sep arrived in Santiago.
    - Reception of line cards (Flexigrid transponders): end of September.
    - Spare parts: Middle of October.
- Santiago–La Serena: This segment was accepted. Telefonica completed the correction of the last elements in the housing sites (final electrical connections).
- La Serena–AURA Gatehouse: This segment is accepted.

02C.08.03.02 – International Chile–US WAN

- *100 Gbps Managed Ring:* No planned engineering activity in this period. The ring is performing as expected.
- *Management and Coordination Contract:* The LSST Network Engineering Team (NET) meeting was held on August 24th. A new confluence page with LSST Network Traffic Types has been created along with the new version of Bandwidth Allocation table, which will focus the future discussion on the minimum allocation only for Science Data Transfer. The Virtual NOC requirements and the 5 seconds requirement for image transfer from La Serena to NCSA were discussed as well. An update on the draft plan for the Network Design and End-to-End Test Plan and QoS planning was briefly presented by Jeronimo Bezerra and Ron Lambert and will be distributed by the next meeting for comments by the NET team. The first Fiber optic light event may not be possible in time for the ADASS Conference due to delayed equipment delivery at REUNA Chile. Dr. Heidi Morgan and Jeronimo Bezerra presented "Large Synoptic Survey Telescope (LSST) Scaling Issues and Network Needs" (part of session Towards A Global Research Platform) at the First National Research Platform (NRP) Workshop, Bozeman, MT, August 7-8. Planning continued for the next South American Astronomy Coordination Committee (SAACC) which will be in Santiago, Chile in conjunction with ADASS conference October 19, 2017. A preliminary program is on the AmLight webpage draft for the SAACC Meeting: https://www.amlight.net/?p=2826.
- *Spectrum Contract:* Angola Cables provided the answers to the technical engineering document detailed in the May 2017 monthly report, which provided the AmLight engineers with the information they needed to begin working to adjust the DWDM system specification accordingly.
- *US National WAN:* ESNet collaboration negotiations with DOE and with FNAL on an operations role continued. FIU continued working with FLR and other fiber providers to determine the most cost effective and greatest throughput for transit from AMPATH north.

Planned activities

02C.08.00 – International Communications and Base Site Management Engineering and Integration

- We expect to receive the Cisco Summit networking equipment in September.

02C.08.01 – Base Center

- Summit to Base ITC Tiger Team: LSE-299 Summit Computer Room Requirements Document to be approved by CCB. New LCR to be prepared to update LSE-309 Summit to Base ITC Design Document with Clean Room and White Room information.
- Summit and Base Networking and Computing: Ron Lambert will visit Tucson to make detailed summit network cable raceway drawings with Bill Schoening, and attend a meeting with Cisco. Complete installation of new firewalls in La Serena.

02C.08.03 – Long-Haul Networks

02C.08.03.01 – Chile National WAN

- Summits - AURA Gatehouse Network: No activity. This segment is accepted.
- DWDM Equipment: We expect to have the Cerro Pachon Caseta new backup generator procured. We expect to receive REUNA DWDM equipment from Coriant (except Spares, which are not necessary to start installation) and start installation. REUNA will attend training in the following courses: "TNMS Operation and Administration for hiT7300" 25-27/Sep and "Tailored TransNet training for hiT 7300" 28-29/Sep.
- Santiago-La Serena: No activity. This segment is accepted.
- La Serena - AURA Gatehouse: No activity. This segment is accepted.

02C.08.03.02 – International Chile - US WAN

- 100 Gbps Managed Ring: AmLight is working to open the 100G circuit in Chile, and we are targeting the ADASS conference to even demonstrate this new capacity. At Level3 Data Center, cross-connects were ordered before the end of August and we are waiting for their delivery. Regarding active equipment, three possible scenarios were created and sent to AURA. Each scenario was described with pros and cons, including different prices.
- Management and Coordination Contract: Continue to coordinate NET meeting and plan for SAACC meeting. Attend the NSF Large Facilities Cyberinfrastructure Workshop in Alexandria VA September 6 - 7. Complete accepted submissions for the ADASS Conference, which will take place in Santiago, Chile in October 22-26, 2017: International Networking in support of Extremely Large Astronomical Data-centric Operations; Challenges of Standardizing and Supporting ITC Services in

a Widely Distributed Project: ITC Design of the LSST Summit - Base Complex.
- Spectrum Contract: The Operations and Maintenance Agreement is expected to be completed.
- US National WAN: Discussions with ESNet for LSST within DOE to define the ESNet service and costs will continue during September. FIU will continue working with FLR and other fiber providers to determine the most cost effective and greatest throughput for transit from AMPATH north.

Recruitment update

- Andres Villalobos accepted the IT System Admin/Eng position offer and will start on September 14. We have completed the hiring for all open Chile ITC positions.

1.02C.10: Science Quality and Reliability Engineering

Current accomplishments

Highlights of work completed this month include:

- Following a significant contribution from Tim Jenness (Architecture) who adapted the LSST stack to the pytest framework, we added junit reporting to Jenkins builds. This work completed a milestone (DLP-338).
- A major refresh of our Jenkins service, including the introduction of the Blue Ocean dashboard, was completed.
- The initial set of documentation for pipe_base has been completed .
- An additional deployment of our Jupyterlab-based environment was stood up on Google's Container Engine for the purposes of running a stack tutorial parallel with the meeting using material provided by Jim Bosch and Simon Krughoff. This was the first time we have had non-team users on our platform and were pleased to note no major issues.
- The migration of the SQuaSH infrastructure to kubernetes is continuing; two more microservices are now complete (the API server and the DB server).
- The team participated at the All Hands meeting. JS talked about Documentation Engineering in the "What is DM working on?" session and did a demo for the "Engaging your Audience" session. FE presented some of the JupyterLab work in the Science Platform session. AT deployed a JupyterLab tutorial cluster (see above).
- AT attended the first ever JupyterCon and gave a summary to the team. Some of our work was mentioned in talks, and we established that the conference is both of strong interest to us and that our work is at the level where we should seriously consider presenting next year.

Planned activities

- FE to visit LCO.

- Mid-cycle checkpoint is underway to schedule work for Sep-Nov.

Pittsburgh contract

- Work performed under this contract included: DM-4953, DM-11410, DM-10521, DM-11463.

Recruitment update

In a highly-anticipated development, Simon Krughoff joined SQuaRE as Scientific Lead on August 1st.

No further recruitment is planned at current budget levels.